

Quine’s “New Foundations” is consistent

M. Randall Holmes

1/2/2017 3:30 PM: version to be submitted

Contents

1	Introduction (Abstract)	2
2	The simple theory of types	3
3	The definition of New Foundations	6
4	Well-known results about New Foundations	8
5	Consistency of NFU	10
6	Tangled type theories	13
6.1	ω - and α -models from tangled type theory	15
7	Tangled webs of cardinals	17
8	The main construction	19
9	Conclusions and questions	53
10	References and Index	55

1 Introduction (Abstract)

In this paper we will present a proof of the consistency of Quine's set theory "New Foundations" (hereinafter NF), so-called after the title of the 1937 paper [11] in which it was introduced. The strategy is to present a Fraenkel-Mostowski construction of a model of an extension of Zermelo set theory without choice whose consistency was shown to entail consistency of NF in our paper [5] of 1995. There is no need to refer to [5]: this paper presents a full (we think a better) account of considerations drawn from that paper.

2 The simple theory of types

New Foundations is introduced as a modification of a simple typed theory of sets which we will call “the simple theory of types” and abbreviate TST (following the usage of Thomas Forster and others).

Definition (the theory TST): TST is a first-order many-sorted theory with equality and membership as primitive relations and with sorts (referred to traditionally as “types”) indexed by the natural numbers. A variable may be written in the form x^i to indicate that it has type i but this is not required; in any event each variable x has a natural number $\text{type}(x)$. In each atomic formula $x = y$, the types of x and y will be equal; in each atomic formula $x \in y$ the type of y will be the successor of the type of x .

The axioms of TST are axioms of extensionality

$$(\forall xy. x = y \leftrightarrow (\forall z. z \in x \leftrightarrow z \in y))$$

for any variables x, y, z of appropriate types and axioms of comprehension, the universal closures of all formulas of the form

$$(\exists A. (\forall x. x \in A \leftrightarrow \phi))$$

for any variables x, A of appropriate types and formula ϕ in which the variable A does not occur.

Definition (set abstract notation): We define $\{x : \phi\}$ as the witness (unique by extensionality) to the truth of the comprehension axiom

$$(\exists A. (\forall x. x \in A \leftrightarrow \phi)).$$

For purposes of syntax, the type of $\{x : \phi\}$ is the successor of the type of x (we allow $\{x : \phi\}$ to appear in contexts (other than binders) where variables of the same type may appear).

This completes the definition of TST. The resemblance to naive set theory is not an accident. This theory results by simplification of the type theory of the famous [20] of Russell and Whitehead in two steps. The predicativist scruples of [20] must first be abandoned, following Ramsey’s [12]. Then it needs to be observed that the ordered pair can be defined as a set, a fact

not known to Whitehead and Russell, first revealed by Wiener in 1914 ([21]). Because Whitehead and Russell did not have a definition of the ordered pair as a set, the system of [20] has a far more complicated type system inhabited by arbitrarily heterogeneous types of n -ary relations. The explicit presentation of this simple theory only happens rather late (about 1930): Wang gives a nice discussion of the history in [19].

The semantics of TST are straightforward (at least, the natural semantics are). Type 0 may be thought of as a collection of individuals. Type 1 is inhabited by sets of individuals, type 2 by sets of sets of individuals, and in general type $n + 1$ is inhabited by sets of type n objects. We do not call the type 0 individuals “atoms”: an atom is an object with no elements, and we do not discuss what elements individuals may or may not have.

Definition (natural model of TST): A natural model of TST is determined by a sequence of sets X_i indexed by natural numbers i and bijective maps $f_i : X_{i+1} \rightarrow \mathcal{P}(X_i)$. Notice that the f_i ’s witness the fact that $|X_i| = |\mathcal{P}^i(X_0)|$ for each natural number i . The interpretation of a sentence in the language of TST in a natural model is obtained by replacing each variable of type i with a variable restricted to X_i (bounding quantifiers binding variables of type i appropriately), leaving atomic formulas $x^i = y^i$ unmodified and changing $x^i \in y^{i+1}$ to $x^i \in f_i(y^{i+1})$.

When we say *the* natural model of TST with base set X_0 , we will be referring to the obvious natural model in which each f_i is the identity map on $X_{i+1} = \mathcal{P}^{i+1}(X_0)$. We may refer to these as default natural models.

Observations about natural models: It is straightforward to establish that

1. The axioms translate to true sentences in any natural model.
2. The first-order theory of any natural model is completely determined by the cardinality of X_0 . It is straightforward to construct an isomorphism between natural models with base types of the same size.

It is usual to adjoin axioms of Infinity and Choice to TST. We do not do this here, and the precise form of such axioms does not concern us at the moment.

The theory TST_n is defined in the same way as TST, except that the indices of the sorts are restricted to $\{0, \dots, n-1\}$. A natural model of TST_n is defined in the obvious way as a substructure of a natural model of TST.

The interesting theory TNT (the “theory of negative types”) proposed by Hao Wang is defined as TST except that the sorts are indexed by the integers. TNT is readily shown to be consistent (any proof of a contradiction in TNT could be transformed to a proof of a contradiction in TST by raising types) and can be shown to have no natural models.

We define a variant TST_λ with types indexed by more general ordinals.

Parameter of the construction introduced: We fix a limit ordinal λ for the rest of the paper.

Definition (type index): A type index is defined as an ordinal less than λ . For purposes of the basic result $\text{Con}(\text{NF})$, $\lambda = \omega$ suffices, but for more general conclusions having more type indices available is useful.

Definition (the theory TST_λ): TST_λ is a first-order many-sorted theory with equality and membership as primitive relations and with sorts (referred to traditionally as “types”) indexed by the type indices. A variable may be written in the form x^i to indicate that it has type i but this is not required; in any event each variable x has an associated $\text{type}(x) < \lambda$. In each atomic formula $x = y$, the types of x and y will be equal; in each atomic formula $x \in y$ the type of y will be the successor of the type of x .

The axioms of TST_λ are axioms of extensionality

$$(\forall xy. x = y \leftrightarrow (\forall z. z \in x \leftrightarrow z \in y))$$

for any variables x, y, z of appropriate types and axioms of comprehension, the universal closures of all formulas of the form

$$(\exists A. (\forall x. x \in A \leftrightarrow \phi))$$

for any variables x, A of appropriate types and formula ϕ in which the variable A does not occur.

In TST_λ , the objects of successor types may be thought of as sets, and the objects of limit types as individuals of various types. Of course, there is not really any interest in TST_λ as such without some relationship postulated between types whose indices do not have finite difference.

3 The definition of New Foundations

The definition of New Foundations is motivated by a symmetry of TST.

Definition (syntactical type-raising): Define a bijection $x \mapsto x^+$ from variables in general to variables with positive type, with the type of x^+ being the successor of the type of x in all cases. Let ϕ^+ be the result of replacing all variables in ϕ with their images under this operation: ϕ^+ is clearly well-formed if ϕ is.

Observations about syntactical type-raising: If ϕ is an axiom, so is ϕ^+ . If ϕ is a theorem, so is ϕ^+ . If $\{x : \phi\}$ is a set abstract, so is $\{x^+ : \phi^+\}$.

This symmetry suggests that the world of TST resembles a hall of mirrors. Any theorem we can prove about any specific type we can also prove about all higher types; any object we construct as a set abstract in any type has precise analogues in all higher types.

Quine suggested that we should not multiply theorems and entities unnecessarily: he proposed that the types should be identified and so all the analogous theorems and objects at different types should be recognized as being the same. This results in the following definition.

Definition (the theory NF): NF is a first-order unsorted theory with equality and membership as primitive relations. We suppose for formal convenience that the variables of the language of TST are also variables of the language of NF (and are assigned the same natural number (or integer) types – though in the context of NF the types assigned to variables do not indicate that they range over different sorts and do not restrict the ways in which formulas can be constructed).

The axioms of TST are the axiom of extensionality

$$(\forall xy. x = y \leftrightarrow (\forall z. z \in x \leftrightarrow z \in y))$$

and axioms of comprehension, the universal closures of all formulas of the form

$$(\exists A. (\forall x. x \in A \leftrightarrow \phi))$$

for any formula ϕ which is a well-formed formula of the language of TST (this is the only context in which the types of variables play a role) and in which the variable A does not occur.

Definition (set abstract notation): We define $\{x : \phi\}$ (for appropriate formulas ϕ) as the witness (unique by extensionality) to the truth of the comprehension axiom

$$(\exists A.(\forall x.x \in A \leftrightarrow \phi)).$$

This is not the way that the comprehension axiom of NF is usually presented. It could make one uncomfortable to define an axiom scheme for one theory in terms of the language of another. So it is more usual to proceed as follows (if this approach is taken there is no need to associate a natural number or integer type with each variable).

Definition: A formula ϕ of the language of NF is *stratified* iff there is a bijection σ from variables to natural numbers (or integers), referred to as a *stratification* of ϕ , with the property that for each atomic subformula ' $x = y$ ' of ϕ we have $\sigma('x') = \sigma('y')$ and for each atomic subformula ' $x \in y$ ' of ϕ we have $\sigma('x') + 1 = \sigma('y')$.

If we were to make more use of stratifications, we would not always be so careful about use and mention. Notice that a formula being stratified is exactly equivalent to the condition that it can be made a well-formed formula of the language of TST by an injective substitution of variables (if we provide as above that all variables of the language of TST are also variables of the language of NF). Of course, if we use the stratification criterion we do not need to assume that we inherit the variables of TST (and their types) in NF.

Axiom scheme of stratified comprehension: We adopt as axioms all universal closures of formulas

$$(\exists A.(\forall x.x \in A \leftrightarrow \phi))$$

for any stratified formula ϕ in which the variable A does not occur.

We discourage any philosophical weight being placed on the idea of stratification, and we in fact make no use of it whatsoever in this paper. We note that the axiom of stratified comprehension is equivalent to a finite conjunction of its instances, so in fact a finite axiomatization of NF can be given that makes no mention of the concept of type at all. However, the very first thing one would do in such a treatment of NF is prove stratified comprehension as a meta-theorem. The standard reference for such a treatment is [3].

4 Well-known results about New Foundations

We cite some known results about NF.

NF as a foundation of mathematics is at least as powerful as TST, since all reasoning in TST can be mirrored in NF.

NF seems to have acquired a certain philosophical cachet, because it appears to allow the formation of large objects excluded from the familiar set theories (by which we mean Zermelo set theory and ZFC) by the “limitation of size” doctrine which underlies them. The universal set exists in NF. Cardinal numbers can be defined as equivalence classes of sets under equinumerousness. Ordinal numbers can be defined as equivalence classes of well-orderings under similarity. We think that this philosophical cachet is largely illusory.

A consideration which one might take into account at this point is that we have not assumed Infinity. TST without Infinity is weaker than Peano arithmetic. TST with Infinity has the same strength as Zermelo set theory with separation restricted to Δ_0 formulas (Mac Lane set theory), which is a quite respectable level of mathematical strength.

In [16], 1954, Specker proved that the Axiom of Choice is refutable in NF, which has the corollary that Infinity is a theorem of NF, so NF is at least as strong as Mac Lane set theory but with the substantial practical inconvenience for mathematics as usually practiced of refuting Choice. It was this result which cast in sharp relief the problem that a relative consistency proof for NF had never been produced, though the proofs of the known paradoxes do not go through.

A positive result of Specker in [17], 1962, served to give some justification to Quine’s intuition in defining the theory, and indicated a path to take toward a relative consistency proof.

Definition (ambiguity scheme): We define the Ambiguity Scheme for TST (and some other similar theories) as the collection of sentences of the form $\phi \leftrightarrow \phi^+$.

Theorem (Specker): The following assertions are equivalent:

1. NF is consistent.
2. TST + Ambiguity is consistent

3. There is a model of TST with a “type shifting endomorphism”, that is, a map which sends each type i bijectively to type $i + 1$ and commutes with the equality and membership relations of the model (it is also equivalent to assert that there is a model of TNT with a type shifting automorphism).

The equivalence also applies to any extension of TST which is closed as a set of formulas under syntactical type-raising and the corresponding extension of NF, and to other theories similar to NF (such as the theories TSTU and NFU described in the next section).

5 Consistency of NFU

In [10], 1969, Jensen produced a very substantial positive result which entirely justified Quine's proposal of NF as an approach to foundations of mathematics, with a slight adjustment of detail.

Define TSTU as a theory with almost the same language as TST (it is convenient though not strictly necessary to add a primitive constant \emptyset^{i+1} in each positive type with the additional axioms $(\forall x.x \notin \emptyset^{i+1})$ for x of each type i) with the same comprehension scheme as TST and with extensionality weakened to allow atoms in each positive type:

Axiom (weak extensionality, for TSTU):

$$(\forall xyz.z \in x \rightarrow (x = y \leftrightarrow (\forall w.w \in x \leftrightarrow w \in y)))$$

Definition (sethood, set abstracts for TSTU): Define $\text{set}(x)$ (x is a set) as holding iff $x = \emptyset \vee (\exists y.y \in x)$ [we are using polymorphism here: the type index to be applied to \emptyset is to be deduced from the type of x]. Define $\{x : \phi\}$ as the witness to the appropriate comprehension axiom as above, with the qualification that if it has no elements it is to be taken to be \emptyset .

Definition (natural models of TSTU): It is convenient to reverse the direction of the functions f_i . A natural model of TSTU is determined by a sequence of sets X_i indexed by natural numbers and a sequence of injections $f_i : \mathcal{P}(X_i) \rightarrow X_{i+1}$. The interpretation of the language of TSTU in a natural model is as the interpretation of the language of TST in a natural model, except that $x^i \in y^{i+1}$ is interpreted as $(\exists z.x^i \in z \wedge f_i(z) = y^{i+1})$. We interpret \emptyset^{i+1} as $f_i(\emptyset)$. It is straightforward to establish that the interpretations of the axioms of TSTU are true in a natural model of TSTU.

Define NFU as the untyped theory with equality, membership and the empty set as primitive notions and with the axioms of weak extensionality, the scheme of stratified comprehension, and the axiom $(\forall x.x \notin \emptyset)$.

Jensen's proof rests on the curious feature that it is possible to skip types in a natural model of TSTU in a way that we now describe. For generality it is advantageous to first present natural models with types indexed by general ordinals less than λ .

Definition (natural models of TSTU_λ): A natural model of TSTU_λ is determined by a sequence of sets X_i indexed by ordinals $i < \lambda$ and a system of injections $f_{i,j} : \mathcal{P}(X_i) \rightarrow X_j$ for each $i < j < \lambda$. Interpretations of the language of TSTU in a natural model of TSTU_λ are provided with a strictly increasing sequence $\{s_i\}_{i \in \mathbb{N}}$ of type indices as a parameter: they are as the interpretation of the language of TST in a natural model, except that each variable x^i of type i is to be interpreted as a variable x^{s_i} restricted to the set X_{s_i} and a membership formula $x^i \in y^{i+1}$ is interpreted as $(\exists z. x^{s_i} \in z \wedge f_{s_i, s_{i+1}}(z) = y^{s_j})$. It is straightforward to establish that the axioms of TSTU have true interpretations in each such scheme. The special constant \emptyset^{i+1} is interpreted as $f_{s_i, s_{i+1}}(\emptyset)$.

Theorem (Jensen): NFU is consistent.

Proof of theorem: Clearly there are natural models of TSTU_λ for each λ : such models are supported by any sequence X_i indexed by $i < \lambda$ with each X_i at least as large as $\mathcal{P}(X_j)$ for each $j < i$. Fix a natural model. Let Σ be any finite set of sentences of the language of TSTU. Let n be a strict upper bound on the type indices appearing in Σ . Define a partition of $[\lambda]^n$: the compartment into which an n -element subset A of λ is placed is determined by the truth values of the sentences in Σ in the interpretation of TSTU in the given natural model parameterized by any sequence s such that the range of $s \upharpoonright n$ is A (the truth values of interpretations of sentences in Σ are determined entirely by this restriction of s). This partition of $[\lambda]^n$ into no more than $2^{|\Sigma|}$ compartments has an infinite homogeneous set H , by the Ramsey theorem, which includes the range of some strictly increasing sequence h of type indices. The interpretation of TSTU determined by h in the natural model satisfies $\phi \leftrightarrow \phi^+$ for each $\phi \in \Sigma$. Thus every finite subset of the Ambiguity Scheme is consistent with TSTU, whence the entire Ambiguity Scheme is consistent with TSTU, and by the results of Specker (the methods of whose proof apply as well to TSTU and NFU as they do to TST and NF), NFU is consistent.

Corollary: NFU is consistent with Infinity and Choice. It is also consistent with the negation of Infinity.

Proof: If X_0 is infinite, all interpretations in the natural model will satisfy

Infinity. If Choice holds in the metatheory, all interpretations in the natural model will satisfy Choice. If all X_i 's are finite (which is only possible if $\lambda = \omega$) the negation of Infinity will hold in the interpreted theory.

Proof without appealing to Specker outlined: Suppose that \leq is a well-ordering of the union of the X_i 's. Add the relation \leq to the language of TSTU. with the same type rules as identity, and interpret $x^i \leq y^i$ as $x^{s_i} \leq y^{s_i}$ when interpreting TSTU using the sequence parameter s as above. We obtain as above a consistency proof for TSTU + Ambiguity + existence of a primitive well-ordering \leq of each type (which can be mentioned in instances of ambiguity). The relation \leq can be used to define a Hilbert symbol: define $(\theta x : \phi)$ as the \leq -least x such that ϕ , or \emptyset if there is no such x . Now construct a model of TSTU + Ambiguity + primitive well-ordering \leq with the same theory consisting entirely of referents of Hilbert symbols (a term model). The Ambiguity Scheme justifies abandoning the distinction between a Hilbert symbol $(\theta x : \phi)$ and its type-raised version $(\theta x^+ : \phi^+)$ in all cases and one obtains a model of NFU with a primitive well-ordering \leq .

The consistency proof for NFU assures us that the usual paradoxes of set theory are indeed successfully avoided by NF, because NFU avoids them in exactly the same ways. This does not rule out NF falling prey to some other unsuspected paradox. Further, though this is not our business here, the consistency proof for NFU shows that NFU is a reasonable foundation for mathematics: NFU + Infinity + Choice is a reasonably fluent mathematical system with enough strength to handle almost all mathematics outside of technical set theory, and extensions of NFU with greater consistency strength are readily obtained from natural models of TST_λ for larger ordinals λ (see [7]).

6 Tangled type theories

In [5], 1995, we pointed out that the method of proof of Jensen can be adapted to NF, establishing the equiconsistency of NF with a certain type theory. This does not immediately give a relative consistency proof for NF, because the type theory under consideration is very strange, and not obviously consistent.

Definition (the theory TTT_λ): TTT_λ (tangled type theory with λ types) is a first-order many-sorted theory with equality and membership as primitive relations and with sorts (referred to traditionally as “types”) indexed by the type indices. A variable may be written in the form x^i to indicate that it has type i but this is not required; in any event each variable x has an associated $\text{type}(x) < \lambda$. In each atomic formula $x = y$, the types of x and y will be equal; in each atomic formula $x \in y$ the type of y will be strictly greater than the type of x .

Let s be a strictly increasing sequence of type indices. Provide a map $(x \mapsto x^s)$ whose domain is the set of variables of the language of TST and whose restriction to type i variables x is a bijection from the collection of type i variables in the language of TST to the collection of type s_i variables in the language of TTT_λ . For each formula ϕ in the language of TST, define ϕ^s as the result of replacing each variable x in ϕ with x^s . We observe that ϕ^s will be a formula of the language of TTT_λ .

The axioms of TTT_λ are exactly the formulas ϕ^s with s any strictly increasing sequence of type indices and ϕ any axiom of TST.

Theorem: TTT_λ is consistent iff NF is consistent.

Proof of theorem: If NF is consistent, one gets a model of TTT_λ by using the model of NF (or if one prefers, disjoint copies of the model of NF indexed by λ) to implement each type, and defining the membership relations of the model in the obvious way.

Suppose that TTT_λ is consistent. Fix a model of TTT_λ . Let Σ be a finite set of sentences of the language of TST. Let n be a strict upper bound on the (natural number) type indices appearing in Σ . We define a partition of $[\lambda]^n$: the compartment into which an n -element subset A of λ is placed is determined by the truth values of the sentences ϕ^s for $\phi \in \Sigma$ for strictly increasing sequences s of type indices such that

the range of $s[n]$ is A . This partition of $[\lambda]^n$ into no more than $2^{|\Sigma|}$ compartments has an infinite homogeneous set H which includes the range of a strictly increasing sequence h of type indices. The interpretation of TST obtained by assigning to each formula ϕ of the language of TST the truth value of ϕ^h in our model of tangled type theory will satisfy each instance $\phi \leftrightarrow \phi^+$ of Ambiguity for $\phi \in \Sigma$. It follows by compactness that the Ambiguity Scheme is consistent with TST, and so by the results of Specker that NF is consistent.

Proof without appealing to Specker outlined: Suppose that \leq is a well-ordering of the union of the types of our model of tangled type theory (this is not an internal relation of the model in any sense). Add the relation symbol \leq to the language of TST. with the same type rules as identity, and transform atomic formulas $x \leq y$ to $x^s \leq y^s$ in the construction of formulas ϕ^s . We obtain as above, using this extended language to define our partition, a consistency proof for TST + Ambiguity + existence of a primitive relation \leq on each type (which can be mentioned in instances of ambiguity, but which cannot be mentioned in instances of comprehension) which is a linear order and a well-ordering in a suitable external sense (any definable nonempty class has a \leq -least element) and which can be used to define a Hilbert symbol: we can define $(\theta x : \phi)$ as the \leq -least x such that ϕ , or \emptyset if there is no such x . Now construct a model of TST + Ambiguity + primitive “external well-ordering” \leq with the same theory consisting entirely of referents of Hilbert symbols (a term model). The Ambiguity Scheme justifies abandoning the distinction between a Hilbert symbol $(\theta x : \phi)$ and its type-raised version $(\theta x^+ : \phi^+)$ in all cases and one obtains a model of NF with a primitive external order \leq .

Examination of our presentation of Jensen’s consistency proof for NFU should reveal that this is an adaptation of the same method to the case of NF. In fact, our “natural model of TSTU $_\lambda$ ” above can readily be understood as a model of TTTU $_\lambda$.

It should also be clear that TTT $_\lambda$ is an extremely strange theory. We cannot possibly construct a “natural model” of this theory, as each type is apparently intended to implement a “power set” of each lower type, and Cantor’s theorem precludes these being honest power sets.

6.1 ω - and α -models from tangled type theory

It is worth noting that the proof of the main result of this paper does not depend on this section: this section is included to indicate why the main result implies further that there is an ω -model of NF.

Jensen continued in his original paper [10] by showing that for any ordinal α there is an α -model of NFU. We show that under suitable conditions on the size of λ and the existence of sets in a model of TTT_λ , this argument can be reproduced for NF.

We quote the form of the Erdős-Rado theorem that Jensen uses: Let δ be an uncountable cardinal number such that $2^\beta < \delta$ for $\beta < \delta$ (i.e., a strong limit cardinal). Then for each pair of cardinals $\beta, \mu < \delta$ and for each $n > 1$ there exists a $\gamma < \delta$ such that for any partition $f : [\gamma]^n \rightarrow \mu$ there is a set X of size β such that f is constant on $[X]^n$ (X is a homogeneous set for the partition of size β).

Let λ be a strong limit cardinal with cofinality greater than $2^{|\alpha|}$. Our types in TTT will be indexed by ordinals $< \lambda$ as usual. We make this stipulation about λ only for this subsection.

We assume the existence of a model of TTT_λ in which each type contains a well-ordering of type α (from the standpoint of the metatheory as well as internally): our language will include names \leq^β for the well-ordering on objects of each type $\beta < \lambda$ and names $[\leq^\beta]_\gamma$ for each $\gamma < \alpha$ for the object of type β at position γ in the order \leq^β . We adjoin similar symbols \leq^i for $i \in \mathbb{N}$ and $[\leq^i]_\gamma$ for $\gamma < \alpha$ to the language of TST. We stipulate that in the construction of formulas ϕ^s , notations \leq^i and $[\leq^i]_\gamma$ will be replaced by notations $\leq^{s(i)}$ and $[\leq^{s(i)}]_\gamma$. A model of NF can then be constructed following the methods above in which a single relation \leq with the objects $[\leq]_\gamma$ in its domain appears. However, the model of NF is obtained by an application of compactness: the order \leq obtained may not be an order of type α or a well ordering at all from the standpoint of the metatheory, because it may have many nonstandard elements. To avoid this, we need to be more careful.

Let Σ_n be the collection of all sentences of the language of TST_n extended as indicated above which begin with an existential quantifier restricted to the domain of an order \leq^i . Let the partition determined by Σ_n make use not of the truth values of the formulas in Σ_n , but of the indices $\gamma < \alpha$ of the minimally indexed witnesses $[\leq]_\gamma$ to the truth of each formula, or α if they are false. The Erdős-Rado Theorem in the form cited tells us that we can find homogeneous sets of any desired size less than λ for this

partition, and moreover (because of the cofinality of λ) we can find, for some fixed assignment of witnesses to each sentence of Σ_n which is witnessed, homogeneous sets of any desired size which induce the fixed assignment of witnesses in the obvious sense. Note that each Σ_n is of cardinality $|\alpha|$ and there are $2^{|\alpha|}$ possible assignments of a witness $\leq \alpha$ to each sentence in Σ_n (recalling that α signals the absence of a witness). This allows us to see that ambiguity of Σ_n is consistent, and moreover consistent with standard values for witnesses to each of the formulas in Σ . We can then extend the determination of truth values and witnesses as many times as desired, because if we expand the set of formulas considered to Σ_{n+1} and partition $(n+1)$ -element sets of type indices instead of n -element sets, we can restrict our attention to a large enough set of type indices homogeneous for the previously given partition (and associated with fixed witnesses) to ensure that we can restrict it to get homogeneity for the partition determined by the larger set of formulas (and get an assignment of witnesses which occurs in arbitrarily large homogeneous sets, as at the previous stage). After we carry out this process for each n , we obtain a full description of a model of TST + Ambiguity with standard witnesses for each existentially quantified statement over the domain of a special well-ordering of type α . We can reproduce our Hilbert symbol trick (add a predicate representing a well-ordering of our model of TTT to the language as above) to pass to a model of NF with the same characteristics.

7 Tangled webs of cardinals

In this section, we replace consideration of the weird type theory TTT_λ with consideration of a (still weird) extension of ordinary set theory (Mac Lane set theory, Zermelo set theory or ZFC) whose consistency is shown to imply the consistency of NF. We are working in set theory without Choice. We note without going into details that we will use Scott's definition of cardinal number, which works for cardinalities of non-well-orderable sets. There is a chain of reasoning in tangled type theory which motivates the details of this definition, but it is better not to present any reasoning in tangled type theory if it can be avoided.

Definition (extended type index, operations on extended type indices):

We define an extended type index as a nonempty finite subset of λ . For any extended type index A , we define A_0 as A , A_1 as $A \setminus \{\min(A)\}$ and A_{n+1} as $(A_n)_1$ when this is defined.

Definition (tangled web of cardinals): A tangled web of cardinals is a function τ from extended type indices to cardinals with the following properties:

naturality: For each A with $|A| \geq 2$, $2^{\tau(A)} = \tau(A_1)$.

elementarity: For each A with $|A| > n$, the first-order theory of the natural model of TST_n with base type $\tau(A)$ is completely determined by the set $A \setminus A_n$ of the smallest n elements of A .¹

Theorem: If there is a tangled web of cardinals, NF is consistent.

Proof of theorem: Suppose that we are given a tangled web of cardinals τ . Let Σ be a finite set of sentences of the language of TST. Let n be a strict upper bound on the natural number type indices appearing in Σ . Define a partition of $[\lambda]^n$: the compartment in which an n -element set A is placed is determined by the truth values of the sentences in Σ in the natural model of TST_n with base type of size

¹We could equally well make the condition $|A| \geq n$, which would simplify the proof immediately below noticeably, but this would not reflect the situation in the actual construction of a tangled web in the next section.

$\tau(A \cup \{\max(A) + 1\})$.² This partition of $[\lambda]^n$ into no more than $2^{|\Sigma|}$ compartments has a homogeneous set H of size $n + 2$. The natural model of TST with base type $\tau(H)$ satisfies all instances $\phi \leftrightarrow \phi^+$ of Ambiguity for $\phi \in \Sigma$: type 1 of this model is of size $2^{\tau(H)} = \tau(H_1)$ and the theory of the natural models of TST with base type $\tau(H_1)$ decides the sentences in Σ in the same way that the theory of the natural models of $\tau(H)$ decides them by homogeneity of H for the indicated partition and the fact that the first order theory of a model of TST_n with base type of size any $\tau(B)$ with $|B| > n$ is determined by the smallest n elements of B . Thus any finite subset of the Ambiguity Scheme is consistent with TST, so $\text{TST} + \text{Ambiguity}$ is consistent by compactness, so NF is consistent by the results of Specker.

There is no converse result: NF does not directly prove the existence of tangled webs. An ω -model of NF will contain arbitrarily large concrete finite fragments of tangled webs.

The existence of a tangled web is inconsistent with Choice. It should be evident that if Choice held in the ambient set theory in which the tangled web is constructed, Choice would hold in the model of NF constructed by this procedure, which is impossible by the results of Specker.

²Having to add one element artificially to A here reflects our decision to require $|A| > n$ in the elementarity condition, which is dictated by the characteristics of the actual construction of a tangled web in the next section.

8 The main construction

The working set theory of the construction is ZFA (with choice). We will carry out a Fraenkel-Mostowski construction to obtain a class model of ZFA (without choice) in which there is a tangled web of cardinals.

cardinal parameters of the construction: We continue to use the previously fixed limit ordinal λ as a parameter. Define A_1 as $A \setminus \{\min(A)\}$ where A is a nonempty finite subset of λ . Define A_0 as A and A_{n+1} as $(A_n)_1$ where this is defined. Define $B \ll A$ as holding iff A and B are distinct, $A \subseteq B$ and all elements of $B \setminus A$ are less than all elements of A ; $B \ll A$ means $B \ll A \vee B = A$. We refer to finite subsets of λ as *clan indices* for reasons which will become evident.

Let κ be a regular uncountable ordinal, fixed for the rest of the paper. We refer to all sets of cardinality $< \kappa$ as *small* and all other sets as *large*.

Let μ be a strong limit cardinal, fixed for the rest of the paper, such that $\mu > |\lambda|$, $\mu > \kappa$, and the cofinality of μ is $\geq \kappa$.

the extent of the atoms described: For each finite subset A of λ , provide a collection of atoms $\mathbf{clan}[A]$ of cardinality μ . If $A \neq B$, $\mathbf{clan}[A]$ and $\mathbf{clan}[B]$ are disjoint. We call these sets *clans*. We provide two other collections of atoms of size μ , disjoint from the clans and from each other, called $\mathbf{parents}[\emptyset]$ and \mathbf{junk} . The elements of the clans will be called *regular atoms* and the elements of the other two sets will be called *irregular atoms*.

our strategy for obtaining a tangled web indicated in advance: The aim of the FM construction is to create the following situation. We use the notation $\mathcal{P}_*(X)$ for the collection of symmetric subsets of a hereditarily symmetric set X . [We haven't said yet what the group and filter are which define the FM construction – we will explain this in due course.]

We consider the default natural models in the FM interpretation with base types $\mathbf{clan}[A]$. The cardinality of $\mathcal{P}_*^{n+2}(\mathbf{clan}[A])$ in terms of the FM interpretation is intended to be determined by A_n where $|A| > n$: $|\mathcal{P}_*^{n+2}(\mathbf{clan}[A])| = |\mathcal{P}_*^2(\mathbf{clan}[A_n])|$ will hold in the FM interpretation. (The clans will be seen directly to be hereditarily symmetric).

The (default) natural models of TST_n in the FM interpretation with bottom types $\mathcal{P}_*^2(\text{clan}[A])$ and $\mathcal{P}_*^2(\text{clan}[B])$ respectively are intended to be isomorphic in the ground ZFA (not in the FM interpretation: the maps witnessing this fact will not be symmetric) if $A \setminus A_n = B \setminus B_n$ (where $|A|, |B| > n$). This implies that they have the same first-order theory as models of TST_n .

Note that if these conditions are achieved, the map τ sending nonempty A to the FM interpretation's cardinality of $\mathcal{P}^2(\text{clan}[A])$ is a tangled web (in the FM interpretation) and the consistency of NF is established.

These effects are to be achieved by careful design of the permutation group and filter generating the FM interpretation.

litter, near-litters, and local cardinals introduced: For each of the clans $\text{clan}[A]$ select a partition of $\text{clan}[A]$ into sets of size κ . We call this partition (a set of subsets of $\text{clan}[A]$) $\text{litters}[A]$. We call the elements of these partitions *litters*. We call the sets $\text{litters}[A]$ themselves *litter partitions* if we have occasion to allude to them: each litter is an element of a litter partition of a clan.

For each clan index A and for each $L \in \text{litters}[A]$, define $[L]$, the *local cardinal*³ of L , as the collection of subsets of $\text{clan}[A]$ with small symmetric difference from L .

We refer to elements of local cardinals as *near-litters*. The local cardinal $[N]$ of a near-litter N is defined as the local cardinal to which N belongs as an element.

For any near-litter N , there is a unique litter N° such that $N \Delta N^\circ$ is small: we refer to the elements of $N \Delta N^\circ$ as the *anomalies* of N .

the global parent function introduced: We will construct a function Π whose domain is the set of all local cardinals (in all clans). This map will have the property that for each $L \in \text{litters}[A]$, $\pi_2(\Pi([L])) = A$ (among many other properties introduced in the course of its specification below).

We will refer to $\pi_1(\Pi([L]))$ as the *parent* of L (and of $[L]$ and of any other near litter N belonging to $[L] = [N]$). The parent of a litter may

³It will be seen to be true in the FM interpretation that subsets of clans belong to the same local cardinal iff they have the same cardinality.

be a set or atom.

The set $\pi_1 \text{“}\Pi \text{“}(\text{dom}(\Pi) \cap \mathcal{P}^2(\text{clan}[A]))$ of all parents of litters included in $\text{clan}[A]$ will be called the *parent set* of $\text{clan}[A]$ and denoted by $\text{parents}[A]$.

We will construct a sequence of approximations Π_α and Π_α^1 ($\alpha \leq \lambda$) to the eventual global parent function $\Pi = \Pi_\lambda$, and it should be noted that notions defined in terms of Π have preliminary versions relative to these approximations. The domain of each of these approximations Π_α and Π_α^1 will be the set of local cardinals of litters included in clans $\text{clan}[A]$ for which the maximum of A is less than α or A is empty. We will use the notation $\Pi_\alpha^{(1)}$ to refer to Π_α or Π_α^1 when we do not care which of these approximations we are referring to.

In particular, it is useful to note that any Π_α which assigns a parent to a given near-litter will assign the same parent to it that Π does, and any Π_α which assigns a parent set to any given $\text{clan}[A]$ will assign the same parent set to it that Π does. This will be evident when we see basic features of the system of approximations. This is not quite true for the other approximations Π_α^1 , as we will see.

allowable permutations introduced: The action of any permutation ρ on a set of atoms S is extended to all sets X whose transitive closure does not contain any atom not in S by the rule $\rho(X) = \rho \text{“} X$ for any set X .

The permutations we use to define the FM interpretation, which we call *allowable permutations*, are exactly those whose action extended to sets fixes Π (and which fix any atoms not in the transitive closure of Π). The allowable permutations in terms of any approximation $\Pi_\alpha^{(1)}$ are similarly those which fix $\Pi_\alpha^{(1)}$ (and atoms not in its transitive closure). We are generally indifferent to what an allowable permutation does outside the transitive closure of the relevant parent function, but we declare that it fixes those atoms for precision.

It is straightforward to see that the allowable permutations must fix clans and fix parent sets (where the clans and parent sets are included in the transitive closure of the parent function used). Suppose that ρ is a permutation which fixes $\Pi_\alpha^{(1)}$ and $\text{clan}[A]$ is a clan included in the transitive closure of $\Pi_\alpha^{(1)}$. Suppose that L is a litter included in

$\mathbf{clan}[A]$. We know that $\pi_2(\Pi_\alpha^{(1)}([L]) = A$, from which it follows that $\pi_2(\Pi_\alpha^{(1)}(\rho([L])) = \rho(A) = A$ (because A is a pure set), so $\rho([L]) \subseteq \mathcal{P}(\mathbf{clan}[A])$. From this it is clear that $\rho(\mathbf{clan}[A]) \subseteq \mathbf{clan}[A]$, and the same argument applied to ρ^{-1} shows that $\rho(\mathbf{clan}[A]) = \mathbf{clan}[A]$. The argument showing that $\mathbf{parents}[A]$ is fixed by ρ is very similar.

It is useful to note that an allowable permutation ρ sends any litter L with parent p to a near-litter N with parent $\rho(p)$ (not necessarily to the *litter* with parent $\rho(p)$: a small collection of atoms may be mapped into or out of N from unexpected litters).

varieties of atoms: We recall that the atoms in the clans are called regular atoms. We stipulate that all irregular atoms which occur in the transitive closure of Π (or of any $\Pi_\alpha^{(1)}$) are parents of litters and do not occur in the transitive closures of any parents of litters which are sets (we will see that this stipulation holds in all the approximations to Π that we consider when we give concrete descriptions of them).

support sets, supports of objects, symmetry: We define a *support set* as a small set of atoms and near-litters in which distinct near-litter elements are disjoint. We say that an object x *has support* S iff every allowable permutation ρ such that $(\forall s \in S : \rho(s) = s)$ also satisfies $\rho(x) = x$. We say that an object is *symmetric* iff it has a support. We say that an object is *hereditarily symmetric* iff it is symmetric and either it is an atom or all elements of its transitive closure are symmetric. Note that it is obvious that any object with a support has a support in which all near-litter elements are actually litters: if S is a support set, the set S° consisting of all atoms in S , all litters N° for $N \in S$, and all elements of $N \Delta N^\circ$ for $N \in S$ is a support set and is a support for any object for which S is a support.

For any hereditarily symmetric set X , we define $\mathcal{P}_*(X)$ as the set of all hereditarily symmetric subsets of X , which we may call the *symmetric power set* of X (relative to a $\Pi_\alpha^{(1)}$ understood in the context).

Note that regular atoms are symmetric with support their own singleton, and irregular atoms are symmetric with support the singleton of any litter of which they are parent.

By standard considerations, the hereditarily symmetric sets and the atoms make up a class model of ZFA, which we will refer to as “the

FM interpretation”, while referring to the ambient ZFA in which we started as “the ground interpretation”.

Of course these notions depend on a particular choice of parent function $\Pi_\alpha^{(1)}$ as a parameter. A support in terms of $\Pi_\alpha^{(1)}$ will have all its elements atoms and near-litters in the transitive closure of $\Pi_\alpha^{(1)}$.

FM interpretations in general: We insert a summary of considerations about FM interpretations in general, taken (and slightly adapted) from [9]:

Let G be a group of permutations of the atoms. Let Γ be a nonempty subset of the collection of subgroups of G with the following properties:

1. The subset Γ contains all subgroups J of G such that for some $H \in \Gamma$, $H \subseteq J$.
2. The subset Γ includes all subgroups $\bigcap C$ of G where $C \subseteq \Gamma$ and C is small [smallness being defined in terms of the parameter κ already introduced above].
3. For each $H \in \Gamma$ and each $\pi \in G$, it is also the case that $\pi H \pi^{-1} \in \Gamma$.
4. For each atom a , $\text{fix}_G(a) \in \Gamma$, where $\text{fix}_G(a)$ is the set of elements of G which fix a .

A nonempty Γ satisfying the first three conditions is what is called a κ -complete normal filter on G .

We call a set A Γ -symmetric iff the group of permutations in G fixing A belongs to Γ . The major theorem which we use but do not prove here is the assertion that the class of hereditarily Γ -symmetric objects (including all the atoms) is a class model of ZFA (usually not satisfying Choice). The assumption that the filter is κ -complete is not needed for the theorem (“finite” usually appears instead of “small”), but it does hold in our construction.

Details of our FM interpretation: We let G be the group of allowable permutations.

For each support set S we define G_S as the subgroup of G consisting of permutations which fix each element of S .

We define the filter Γ as the set of subgroups H of G which extend subgroups G_S . The only point which requires special comment in the verification that Γ is a normal filter is the normality condition: it is straightforward to establish that if $\pi \in G$ and $G_S \subseteq H \in \Gamma$, then $G_{\pi(S)} \subseteq \pi H \pi^{-1}$, establishing normality.

Note that an FM interpretation is obtained in this way in terms of each $\Pi_\alpha^{(1)}$.

small subsets of the domain of the FM interpretation: It should be evident that any set of cardinality $< \kappa$ of hereditarily symmetric sets is hereditarily symmetric: choose a support for each element of the small set, and the union of the chosen supports of the elements of the set can serve as a support for the set (mod a technical point: assume that all near-litters in the supports chosen for the elements are litters, so that the union can be relied on to be a support set).

the intended description of parent sets: As already noted, the set $\text{parents}[\emptyset]$ is a set of irregular atoms of size μ .

We intend that for each nonempty A ,

$$\text{parents}[A] = \text{clan}[A_1] \cup \bigcup_{B \ll A} \mathcal{P}_*^{|B|-|A|+1}(\text{clan}[B])$$

will hold.

As contemplation of this formula may suggest to the reader, it takes work to show that the intention just expressed is possible to realize.

The superscript in the notation $\mathcal{P}_*^{|B|-|A|+1}(\text{clan}[B])$ indicates finite iteration of the symmetric power set operation.

construction of the global parent function starts: We will construct Π in stages. For each α , we construct a function which will eventually be seen to be the restriction of Π to local cardinals of litters in clans $\text{clan}[A]$ with A empty or $\max(A) < \alpha$, which we call Π_α . There will be other intermediate approximations Π_α^1 where α is a successor ordinal.

Note that the relationship between Π_α and Π which is projected will ensure that where Π_α assigns parents to near-litters or parent sets to clans, Π will assign the same parents and parent sets.

the zero stage: The case $\alpha = 0$ is straightforward: we fix a bijection from the set of local cardinals of elements of $\mathbf{litters}[\emptyset]$ to $\mathbf{parents}[\emptyset] \times \{\emptyset\}$ as Π_0 : we have already chosen these sets and we know them to be of the same cardinality μ .

inductive hypotheses: At any stage α , we assume that we have already constructed Π_β for each $\beta < \alpha$ and that for each $\beta < \gamma < \alpha$ we have $\Pi_\beta \subseteq \Pi_\gamma$, and that the formulas for parent sets given above hold for each Π_β with $\beta < \alpha$ for clans $\mathbf{clan}[A]$ with the maximum of A less than β (with the symmetric power set \mathcal{P}_* being defined in terms of allowable permutations defined in terms of Π_β instead of the full Π not yet defined).

We have further inductive hypotheses needed to drive the construction: here we state these and verify that they hold at the basis (where there is in fact anything to prove at the basis).

1. For any natural number n , $\mathcal{P}_*^n(\mathbf{clan}[A])$ is the same set as defined in terms of any two Π_β, Π_γ with $\mathbf{max}(A) < \beta < \gamma < \alpha$.
2. For A, B nonempty finite subsets of λ and $0 < n < |A|$ and $0 < n < |B|$, there is an external isomorphism (in the ground interpretation of ZFA rather than the FM interpretation: the function realizing the isomorphism is not symmetric) between the default natural models of \mathbf{TST}_{n+2} in the FM interpretation with base types $\mathbf{clan}[A]$ and $\mathbf{clan}[B]$ and top types $\mathcal{P}_*^{n+1}(\mathbf{clan}[A])$ and $\mathcal{P}_*^{n+1}(\mathbf{clan}[B])$, respectively, if $A \setminus A_n = B \setminus B_n$ – this being asserted in terms of each Π_β with $\beta < \alpha$, for each A, B with maximum less than β .
3. A *locally small bijection* is defined as a bijection whose domain and range are the same set of atoms and which sends any atom in its domain which belongs to a clan to an atom in the same clan and sends irregular atoms to irregular atoms, and whose domain has small intersection with each litter (the domain may contain none, some, or all of the irregular atoms).

We say that a regular atom x belonging to a litter L is an *exception* of an allowable permutation ρ iff $\rho(x) \notin \rho(L)^\circ$. (recall that for any near-litter N , N° is defined as the litter with small symmetric difference from N).

We stipulate as an inductive hypothesis that any locally small bijection ρ_0 extends to an allowable permutation ρ , with no exceptions other than elements of the domain of the locally small bijection ρ_0 , with all concepts being defined in terms of any appropriate Π_β for $\beta < \alpha$ (β must dominate the indices of the clans represented in the domain of ρ_0). We call this condition the *extension property*. We say that such an allowable permutation ρ is a *substitution extension* of the locally small bijection ρ_0 .

We verify that the extension property holds at the basis $\alpha = 0$. Let ρ_0 be a bijection on a set of atoms A included in $\mathbf{parents}[\emptyset] \cup \mathbf{clan}[\emptyset]$, the set A having small intersection with each litter in $\mathbf{clan}[\emptyset]$, with the further conditions that $\rho_0 \text{“}\mathbf{parents}[\emptyset] \subseteq \mathbf{parents}[\emptyset]\text{”}$ and $\rho_0 \text{“}\mathbf{clan}[\emptyset] \subseteq \mathbf{clan}[\emptyset]\text{”}$. Extend ρ_0 (giving the new map the same name) to act as the identity on $\mathbf{parents}[\emptyset] \setminus A$. For each litter L , find the litter M such that $\pi_1(\Pi_0([M])) = \rho_0(\pi_1(\Pi_0([L])))$. Select a bijection ρ_L from $L \setminus \mathbf{dom}(\rho_0)$ to $M \setminus \mathbf{dom}(\rho_0)$. The desired allowable permutation is the union of ρ_0 and all the maps ρ_L : it is clearly allowable and has no exceptions other than elements of the domain of ρ_0 .

4. A *strong support set* is defined as a support set S on which there is a (strict) well-ordering $<_S$ under which any regular atom x in S is preceded in $<_S$ by the near-litter in S which contains x , if there is one (there is at most one such near-litter because distinct near-litter elements of a support set are disjoint), and which satisfies the condition that if the parent p of any near-litter N in S belongs to a set $\mathcal{P}^{n+1}(\mathbf{clan}[B])$, then p has the set

$$\{x : x <_S N \wedge (\exists C \leq B_n : x \in \mathbf{clan}(C) \vee x \in \mathbf{litters}[C])\}$$

as a (strong) support, and if the parent of a near-litter N in S is an atom x , the atom x either precedes N in $<_S$ or does not belong to S [the latter must of course be the case if x is irregular].

We introduce the notation S_γ for the element $s \in S$ such that the restriction of $<_S$ to $\{u : u <_S s\}$ has order type γ .

An *extended strong support set* has the further properties that every near-litter belonging to the support set is a litter and every atom in the support belongs to a litter in the support, and every

regular atom which is the parent of a litter in the support belongs to the support.

We stipulate as an inductive hypothesis that every regular atom has an extended strong support (in terms of Π_β for $\beta < \alpha$). From this it is straightforward to establish that any support set can be extended to a strong support set (and indeed to an extended strong support set), and so every object with a support has an (extended) strong support. It is obvious that each regular atom in the transitive closure of Π_0 has an extended strong support in terms of Π_0 , whose elements are the atom itself and the litter to which it belongs (whose parent is an irregular atom).

5. Each element of $\mathcal{P}_*^{n+1}(\text{clan}[B])$ [for $n \leq |B|$] has a strong support all of whose elements belong to clans $\text{clan}(C)$ and litter sets $\text{litters}(C)$ such that $C \ll B_n$. We call such a support a *relevant support* (or relevant strong support for emphasis). We further stipulate that a relevant support of an element of V_ω (a hereditarily finite pure set) is empty. The motive for this last condition is that hereditarily finite pure sets are the only sets which can belong to more than one set of the form $\mathcal{P}^n(\text{clan}[B])$: we do not want their relevant supports to contain information about them which is tied to a specific set of this form, and since they are invariant there is no reason to allow this.

This obviously holds at the basis, where every support can be supplied with an order that makes it strong and every element of any support is in $\text{clan}[\emptyset]$ or $\text{litters}[\emptyset]$, so every element of $\mathcal{P}(\text{clan}[\emptyset])$ (the only iterated power set to be considered at this stage) has the desired property.

6. This clause does not state a new inductive hypothesis, but points out corollaries of the extension property and the existence of extended strong supports. The general intention of this subsection is to clarify the extent to which allowable permutations act freely on general objects. The extension property tells us that they act quite freely on atoms, but since objects also have near-litters in their supports, whose parents may be sets, the degree of freedom of action on general objects requires clarification.
 - (a) Let S and T be strong supports with orders $<_S$ and $<_T$, respectively. We give a set of conditions equivalent to existence

of an allowable permutation mapping $<_S$ to $<_T$.

- i. For each appropriate γ and clan index C , $S_\gamma \in \mathbf{clan}[C] \leftrightarrow T_\gamma \in \mathbf{clan}[C]$ and $S_\gamma \subseteq \mathbf{clan}[C] \leftrightarrow T_\gamma \subseteq \mathbf{clan}[C]$
- ii. For each appropriate γ, δ , $S_\gamma \in S_\delta$ iff $T_\gamma \in T_\delta$.
- iii. For each appropriate γ, δ , the atom S_γ is the parent of the litter S_δ iff the atom T_γ is the parent of the litter T_δ .
- iv. For each appropriate γ , S_γ has set parent iff T_γ has set parent, and under this condition there is an allowable permutation mapping the initial segment of $<_S$ with order type γ to the initial segment of $<_T$ with order type γ , and any such map sends the parent of S_γ to the parent of T_γ (if any such map does this, all such maps do, because they agree on a support of the parent of S_γ).

That these conditions hold if $<_T$ is the image of $<_S$ under an allowable permutation is evident.

We establish the converse using the extension property.

Let χ be the common order type of $<_S$ and $<_T$.

We indicate how to construct, for each $\gamma \leq \chi$, a locally small bijection ρ_0^γ with suitable properties, by recursion. For $\gamma < \delta \leq \chi$ we will have $\rho_0^\gamma \subseteq \rho_0^\delta$. The inductive hypotheses of the recursion are that for any substitution extension ρ^γ of ρ_0^γ and any $\delta < \gamma$, we have $\rho^\gamma(S_\delta) = T_\delta$.

- i. For $\lambda \leq \chi$ limit, ρ_0^λ is the union of all ρ_0^γ for $\gamma < \lambda$.
- ii. For each γ such that S_γ is an atom or a litter with parent an atom, we will have $\rho_0^{\gamma+1}$ a locally small bijection extending $\rho_0^\gamma \cup \{(S'_\gamma, T'_\gamma)\}$, where S'_γ is defined as S_γ if S_γ is an atom and otherwise as the parent of S_γ , and T'_γ is defined similarly [an at most countably infinite set of additional values may be required to preserve the fact that the map is a bijection with the same set as domain and range: these should not be elements of S or T or parents of elements of S or T , and if such an additional element belongs to an S_δ it must be mapped to an element of T_δ]. It is immediate that $\rho^{\gamma+1}(S_\gamma) = T_\gamma$ for any substitution extension $\rho^{\gamma+1}$ of $\rho_0^{\gamma+1}$.
- iii. For each γ such that S_γ is a near-litter with parent a set, note that a substitution extension ρ^γ of ρ_0^γ will map the

parent of S_γ to the parent of T_γ because it acts correctly on each element of a support of S_γ . We extend ρ_0^γ to $\rho_0^{\gamma+1}$ by the following procedure: associate with each anomaly s of S_γ (not in S or in the domain of ρ_0^γ) a sequence of atoms s_i with $s_0 = s$ and each other s_i in S_γ° and with each anomaly t of T_γ (not in T or in the domain of ρ_0^γ) associate a sequence of atoms t_i with $t_0 = t$ and each other t_i in T_γ° , and further for each anomaly s of S_γ not in S provide a sequence s'_i of atoms in T_γ° and for each anomaly t of T_γ not in T provide a sequence t'_i of atoms in S_γ° , all of these sequences being injective and the ranges of all these sequences being disjoint from each other and from S and T and from the domain of ρ_0^γ : $\rho_0^{\gamma+1}$ is a locally small bijection which extends ρ_0^γ and contains all pairs (s_i, s'_i) and all pairs (t'_i, t_i) [as noted above, the additional values added to preserve the fact that the map is a bijection with the same set as domain and range should not be in S , T , or be parents of elements of S or T and if such an additional element belongs to an S_δ it must be mapped to an element of T_δ ; only a small collection of such values are needed (no more than countably many values per explicitly described value, to fill out orbits in $\rho_0^{\gamma+1}$)]. This causes the desired condition $\rho^{\gamma+1}(S_\gamma) = T_\gamma$ to hold for any substitution extension $\rho^{\gamma+1}$ of $\rho_0^{\gamma+1}$.

Each ρ_0^γ is clearly a locally small bijection, and it should be evident that a substitution extension of ρ_0^χ will map each S_γ to T_γ and so send $<_S$ to $<_T$ as desired.

It is also worth noting that the argument here shows that if S and T are strong supports with orders $<_S$ and $<_T$ of limit order type in which it is possible for each pair of proper initial segments S', T' in the given order on S, T respectively of the same length (with restricted orders $<_{S'}, <_{T'}$) to find an allowable permutation ρ' such that $\rho'(<_{S'}) = \rho(<_{T'})$, we can find an allowable permutation ρ such that $\rho(<_S) = \rho(<_T)$.

- (b) Let ρ_0 be a locally small bijection whose domain includes all irregular atoms. Using a $\Pi_\beta^{(1)}$ for $\beta < \alpha$, large enough to

handle the domain of ρ_0 , choose for each pair of litters L, M which are in the same clan a bijection $\rho_{L,M}$ from $L \setminus \text{dom}(\rho_0)$ to $M \setminus \text{dom}(\rho_0)$. Then ρ_0 extends to a **unique** allowable permutation ρ which also extends $\rho_{L,\rho(L)^\circ}$ for each L .

We prove this by presenting a procedure for computing a function ρ for each regular atom or litter x and then arguing that the action of ρ on atoms must determine the unique allowable permutation with the indicated properties.

We choose a strict well-ordering $<_{\text{a11}}$ of all regular atoms and litters with the property that the weak segment

$$\{y : y <_{\text{a11}} x \vee y = x\}$$

contains an extended strong support of x for each regular atom or litter x (on which the order used agrees with $<_{\text{a11}}$). Such a well-ordering clearly exists. Segments in this order fail to be extended strong supports only by being too large.

We fix a regular atom or litter x and assume that the value of $\rho(y)$ has been computed for each element y of $\{y : y <_{\text{a11}} x\}$. It is also an inductive hypothesis that ρ is not assigned exceptions not in the domain of ρ_0 ; this is readily seen to be preserved by the procedure described below.

If x is a regular atom, then it either belongs to the domain of ρ_0 , in which case $\rho(x)$ is defined as $\rho_0(x)$ [we can regard these assignments of values as made immediately at the outset], or it does not belong to the domain of ρ_0 : in this case it belongs to a litter $L <_{\text{a11}} x$. By hypothesis of the recursion, $\rho(L)$ has been computed already, and we define $\rho(x)$ as $\rho_{L,\rho(L)^\circ}(x)$.

If x is a litter L , then it has a parent p : once we have computed a value $\rho^*(p)$ [how we do this is described just below], we find the litter M with parent $\rho^*(p)$ included in the same clan as L and compute $\rho(L)$ as $\rho_0 \text{“} L \cup \rho_{L,M} \text{”} L$. To avoid mystification, we point out that $\rho^*(p)$ will be seen in the end to be just $\rho(p)$ in the usual sense, but we cannot presume this during the recursion.

We have to describe the computation of $\rho^*(p)$. If p is an irregular atom, we define $\rho^*(p)$ as $\rho_0(p)$. If p is a regular atom, we define $\rho^*(p)$ as $\rho(p)$, which we have by hypothesis of the recursion already computed. If p is a set, it has an extended strong

support S all of whose elements are $<_{\mathbf{a11}} x$, with $<_S$ a restriction of $<_{\mathbf{a11}}$. We have by hypothesis of the recursion already computed ρ at all values in S . We construct a substitution extension ρ' of a locally small bijection extending the union of the restriction of ρ to atoms in S and the map ρ_0 , with a further condition: some additional values $\rho'(z)$ may have to be assigned to preserve the condition that a locally small bijection is a bijection with domain and range the same set: for any litter q in S , we must preserve $z \in q \leftrightarrow \rho'(z) \in \rho(q)$, avoiding creation of additional exceptions of ρ' or images under ρ' of exceptions of ρ' belonging to litters in S . Note that any exception of ρ' which either belongs to a litter in S or is mapped into a litter in S must be a value at which ρ has already been defined. We define $\rho^*(p)$ as $\rho'(p)$. There are things that need to be verified to ensure that this works. We note that each litter q in S satisfies $\rho'(q) = \rho(q)$; if this were not the case, there would be a $<_{\mathbf{a11}}$ -first such q , to which we confine our attention: the parent of q clearly has the same image under ρ' that it does under ρ , since elements of a support for q included in S satisfy this condition: for $\rho(q)$ to be distinct from $\rho'(q)$, one of the following things must happen:

- i. an element r of q must satisfy $\rho'(r) \notin \rho(q)$. If $\rho'(r) \in \rho(q)^\circ$, then $\rho^{-1}(\rho'(r))$ is an exception of ρ , so belongs to the domain of ρ_0 , so r itself is in the domain of ρ_0 , so $\rho(r) = \rho'(r)$, a contradiction. If $\rho'(r) \notin \rho(q)^\circ$, then r is an exception of ρ' , so $\rho(r)$ is already defined so as to agree with $\rho'(r)$, a contradiction.
- ii. an $r \notin q$ must satisfy $\rho'(r) \in \rho(q)$. If $\rho'(r) \in \rho(q)^\circ$, then r is an exception of ρ' , so $\rho(r)$ has already been defined as $\rho'(r)$, a contradiction. If $\rho'(r) \notin \rho(q)^\circ$, then $\rho^{-1}(\rho'(r))$ is an exception of ρ , and so belongs to the domain of ρ_0 , and so r belongs to the domain of ρ_0 , and $\rho(r) = \rho'(r)$ must hold, a contradiction.

This establishes that any particular choice of extended strong support S for x from the $<_{\mathbf{a11}}$ -segment of x determines a unique value for $\rho^*(p)$, because values at all elements of the support are uniquely determined. It cannot be the case that

two distinct choices of extended support give different values of $\rho^*(p)$, because the union of two extended supports is an extended support, and the value computed from the union cannot disagree with either of the values computed from the original distinct extended supports.

The function ρ computed by the procedure above clearly has the action on atoms of an allowable permutation satisfying the indicated conditions. If there were two distinct allowable permutations ρ_1, ρ_2 satisfying the given conditions, there would be a \leq_{all} -first value x at which they disagreed: this value is clearly not an atom or a litter with atomic parent, and almost as clearly not a litter with set parent, by considering the identical action of ρ_1 and ρ_2 on the support of the parent of such a litter and the fact that ρ_1 and ρ_2 have no exceptions at which they can disagree (all exceptions of either map being in the domain of ρ_0).

Note that the value of ρ computed at an object is completely determined by the values of ρ_0 at atoms (regular and irregular) occurring in the given extended strong support of the object (in the case of irregular atoms, occurring as parents of litters in the extended strong support), along with the maps $\rho_{L,M}$ fixed in advance.

limit stages: When α is limit, we define Π_α as $\bigcup_{\beta < \alpha} \Pi_\beta$. The only induction hypotheses which require much attention at the limit stages are induction hypothesis 1 and the extension property, most conveniently discussed below after the verification at successor stages is described.

successor stage introduced: Now suppose that $\alpha = \beta + 1$.

We have already constructed Π_β by the hypothesis of the recursion.

the first approximation Π_α^1 to Π_α : We first construct an initial approximation to Π_α , which we call Π_α^1 .

Choose a bijection σ_β whose domain is the set of all atoms in the transitive closure of Π_β , sending **parents** $[\emptyset]$ to **junk**, and sending each **litters** $[A]$ to **litters** $[\{\beta\} \cup A]$, where β dominates all elements of A (that is, sending **clan** $[A]$ to **clan** $[\{\beta\} \cup A]$ in a way which preserves

the litter structure). Define **fix** on the range of Π_β by $\mathbf{fix}(x, A) = (x, \{\beta\} \cup A)$.

Define Π_α^1 as $\Pi_\beta \cup \sigma_\beta(\mathbf{fix} \circ \Pi_\beta)$, where the action of σ_β on a set whose transitive closure includes no atom not in the domain of σ_β is defined by the rule $\sigma_\beta(A) = \sigma_\beta A$. Notice that Π_α^1 is best understood as consisting of two isomorphic copies of Π_β . It satisfies all the conditions on parent sets of clans except that $\mathbf{parents}[\{\beta\}]$ is a new set **junk** of μ irregular atoms (already provided earlier: we will see that the same set can be used at each successor stage, because **junk** is purged from the transitive closure of Π_α).

It is important to notice that the symmetric sets in any $\mathcal{P}_*^n(\mathbf{clan}[B])$ with $\mathbf{max}(B) < \beta$, with the symmetry defined in the sense of Π_α^1 , are exactly the symmetric sets in $\mathcal{P}_*^n(\mathbf{clan}[B])$, with the symmetry defined in the sense of Π_β , and the symmetric sets in $\mathcal{P}_*^n(\mathbf{clan}[\{\beta\} \cup B])$, with the symmetry defined in the sense of Π_α^1 , are exactly the images under σ_β of symmetric sets in $\mathcal{P}_*^n(\mathbf{clan}[B])$ with the symmetry defined in the sense of Π_β . This holds because the actions possible for allowable permutations on $\mathcal{P}_*^n(\mathbf{clan}[\{\beta\} \cup B])$ are exactly those which are analogous to actions on $\mathcal{P}_*^n(\mathbf{clan}[B])$ via the action of σ_β extended to sets, and also because the action of allowable permutations on clans included in the transitive closure of Π_β and on clans included in the transitive closure of $\sigma_\beta(\mathbf{fix} \circ \Pi_\beta)$ are completely independent of one another, as no atom appearing in the transitive closure of one of these sets appears in the transitive closure of the other. In symbols, the allowable permutations in the sense of Π_α^1 are exactly the permutations $\rho_1 \cup \sigma_\beta \rho_2 \sigma_\beta^{-1}$, where ρ_1 and ρ_2 are any independently chosen allowable permutations in the sense of Π_β [here we regard an allowable permutation in the sense of $\Pi_\gamma^{(1)}$ as determined by its restriction to atoms in the transitive closure of $\Pi_\gamma^{(1)}$].

the construction of Π_α by repairing(?) the parent set of $\mathbf{clan}(\{\beta\})$:

What we do next is repair Π_α^1 so that $\mathbf{clan}[\{\beta\}]$ has the intended parent set

$$\mathbf{clan}[\emptyset] \cup \bigcup_{B \ll \{\beta\}} \mathcal{P}_*^{|B| - |\{\beta\}| + 1}(\mathbf{clan}[B])$$

which we will call Ω_α . The symmetric power set here is defined in terms of Π_α^1 .

The preceding expression simplifies to

$$\Omega_\alpha = \mathbf{clan}[\emptyset] \cup \bigcup_{B \ll \{\beta\}} \mathcal{P}_*^{|B|}(\mathbf{clan}[B])$$

The set component of this is externally isomorphic (via application of σ_β) to

$$\begin{aligned} & \bigcup_{\max(B) < \beta} \mathcal{P}_*^{|B|+1}(\mathbf{clan}[B]) : \\ & \sigma_\beta \left(\bigcup_{\max(B) < \beta} \mathcal{P}_*^{|B|+1}(\mathbf{clan}[B]) \right) \\ &= \bigcup_{\max(B) < \beta} \mathcal{P}_*^{|B|+1}(\mathbf{clan}[B \cup \{\beta\}]) \\ &= \bigcup_{B \ll \{\beta\}} \mathcal{P}_*^{|B|}(\mathbf{clan}[B]), \end{aligned}$$

(the last step being a reindexing, replacing references to B with references to $B \cup \{\beta\}$),

which is entirely describable in terms of the symmetry already computed from Π_β . It should be clear that the definition of this last set is unaffected by the choice of Π_β or Π_α^1 as one's working parent function (approximation to Π).

Our intention is to define Π_α as agreeing with Π_α^1 everywhere except on local cardinals of litters included in $\mathbf{clan}[\{\beta\}]$: these local cardinals will instead be mapped bijectively to Ω_α (which we should recall is defined in terms of symmetric power sets in terms of Π_α^1). For this to work, we need to show that Ω_α is of cardinality no greater than μ (it is clearly of cardinality at least μ). We also need to exercise some care with regard to the details of the bijection from local cardinals of elements of $\mathbf{clan}[\{\beta\}]$ to the complex parent set. Notice that the additional set of atoms **junk** used as **parents** $[\{\beta\}]$ in Π_α^1 disappears completely from

Π_α (so we can actually use the same collection of additional atoms at each successor stage).

To be more exact, let h be a bijection from σ_β “**parents**[\(\emptyset\)] = **junk**” to Ω_α [we have not forgotten that we need to show that these two sets are of the same cardinality μ : this is attended to below] and define $\Pi_\alpha([N])$ as $(h(\pi_1(\Pi_\alpha^1([N])), \{\beta\}))$ if $N \subseteq \mathbf{clan}[\{\beta\}]$ and as $\Pi_\alpha^1([N])$ otherwise.

To arrange for every set to have an extended strong support in the sense of Π_α requires that we pay a little extra attention to how h is chosen. The requirement is that we start with a well-ordering (for which we will have an occasion to use a name $<_h$) of σ_β “**parents**[\(\emptyset\)] = **junk**” (hereinafter “the domain”) and a well-ordering of

$$\Omega_\alpha = \mathbf{clan}[\emptyset] \cup \bigcup_{B \ll \{\beta\}} \mathcal{P}_*^{|B|}(\mathbf{clan}[B])$$

(hereinafter “the range”), both of order type μ , then proceed to define h in steps indexed by the ordinals below μ , at each step choosing the image under h of the first unused element x of the domain to be the first unused element of the range which has an extended strong support which does not include any litter with an irregular atom $z \not\prec_h x$ as parent. We can be sure that h thus defined will include all of its intended range because μ has cofinality $\geq \kappa$.

motivational remark: It is useful to note for motivation that the effect of replacing **junk** with Ω_α as **parents**[\(\{\beta\}\)] is to remove some allowable permutations of **parents**[\(\{\beta\}\)] and so possibly to add sets to each $\mathcal{P}_*^n(\mathbf{clan}[\{\beta\}])$. Further analysis will reveal that the effect is to add sets to $\mathcal{P}_*^n(\mathbf{clan}[B])$ when $\max(B) = \beta$ and $n > |B|$, and not to add any sets to iterated power sets of clans with lower index.

combinatorics of power sets of clans We show that the symmetric power set of any clan is exactly the collection of sets which have small symmetric difference from a small or co-small union of litters.

Suppose X is a symmetric subset of $\mathbf{clan}[A]$. Let S be an extended strong support of X .

Let L be a litter. Suppose $L \cap X$ and $L \setminus X$ are both large. Choose a, b belonging to $L \cap X$ and $L \setminus X$ which are not elements of S . There

is a locally small bijection which swaps a, b , and fixes each atom belonging to S and each irregular atom. A substitution extension of this locally small bijection, that is, an allowable permutation extending the locally small bijection and with no exceptions outside the domain of the locally small bijection, will fix each element of S because if it did not, there would be a first element in the order on S which was moved, and it would be a litter, and its parent would not be moved, so the litter would have to have an exception of the permutation or an image of an exception among its elements, which would have to be a or b , and a, b are not moved out of the litter to which they belong by this permutation. So this map cannot move X because it fixes all elements of its support, but also clearly moves X . This contradiction shows that no symmetric set can cut a litter into two large parts.

Suppose X cuts each of a large collection of litters. Let S be an extended strong support for X . Choose a litter L which does not belong to S and contains no element of S and is cut by X . Choose $a \in L \cap X$ and $b \in L \setminus X$. Consider the locally small bijection interchanging a, b and fixing each atom in S and irregular atom. The argument that an allowable permutation extending this locally small bijection (with no exceptions outside the domain of the locally small bijection) both fixes and does not fix X goes exactly as in the previous paragraph. Thus any symmetric subset X of a clan has small symmetric difference from a union of litters.

Now suppose that X includes the union of a large collection of litters and fails to meet the union of another large collection of litters. Choose litters L included in X and M not meeting X , neither belonging to the extended strong support S of X . Choose $a \in L$ and $b \in M$, neither belonging to S . Now extend the locally small bijection interchanging a and b and fixing each atom in S and irregular atom to an allowable permutation with no exceptions outside the domain of the locally small bijection. Suppose this moves any element of S : the first element in the sense of $<_S$ which is moved must be a litter, its parent must be fixed, so the litter must include an exception or an image of an exception of the map which is moved by the map. But the only exceptions of the map which are moved (and the only images of such exceptions) are a, b , which belong to litters which are not in S . Thus X is fixed. And yet of course the map moves X . This contradiction completes the proof that

any symmetric subset of a clan has small symmetric difference from a small or co-small union of litters.

Further, it is obvious that any subset of a clan with small symmetric difference from a small or co-small union of litters is actually symmetric.

Clan subset support lemma: If S is a strong support of an element Z of $\mathcal{P}(\text{clan}[B])$ then Z is expressible as the symmetric difference of a set $X \subseteq S$ of atoms and the union or the complement of the union of a set $Y \subseteq S$ of near-litters.

Proof of lemma: Let S be a strong support of an element Z of $\mathcal{P}(\text{clan}[B])$, which is the symmetric difference of a small set X' of atoms and either the union or the complement of the union of a small set Y' of litters by results shown above.

For an allowable permutation ρ to fix Z , it is sufficient for ρ to be a substitution extension of a locally small bijection ρ_0 fixing each atomic element of S and atomic parent of an element of S , assigning values to each anomaly of an element of S , and compatible with fixing each near-litter N in S in the sense that for each x in the domain of ρ_0 , $\rho_0(x) \in N \leftrightarrow x \in N$: suppose that such a map moved Z ; consider the $<_S$ -first element u which it moves, which must be a near-litter in S : the parent of this near-litter is fixed because it has a support consisting of things appearing earlier in $<_S$; so some element of u is mapped to a non-element of u or vice versa: neither this element nor its image can be an anomaly of u , so this element must be an exception of ρ , and ρ has no exceptions which it moves in a way not compatible with fixing a near-litter in S .

Suppose that $x \in X' \setminus S$. Let $y \neq x$ belong to any near-litter in S which contains x (there might not be such a near-litter, in which case y is chosen not to belong to any near-litter in S) with $y \notin X'$ and $y \notin S$: a substitution extension of the map fixing all atoms which belong to S or are parents of near-litters in S or are anomalies of near-litters in S (other than x or y if either of them happens to be such an anomaly) and in addition swapping x and y fixes Z by considerations above, but at the same time clearly moves Z (it is worth noting in this connection that x and y either both belong to $\bigcup Y'$ or both do not belong to $\bigcup Y'$); this

contradiction shows that all elements of X' belong to S .

Suppose that $\bigcup Y'$ is not included in the union of a collection $X_1 \subseteq S$ of atoms and the set union of a collection $Y_1 \subseteq S$ of near-litters. This implies that we can find $z \in \bigcup Y'$ which does not belong to S (and so does not belong to X') and belongs to a near-litter which does not belong to S . Find $w \notin \bigcup Y'$ which does not belong to S (and so not to X') nor to any near-litter in S ; a substitution extension of the map which fixes each atom in S , atomic parent of an element of S , anomaly of an element of S (other than z or w if either happens to be such an anomaly), and in addition swaps z and w will again both move and not move Z , so the union of Y' is in fact included in a set $X_1 \cup \bigcup Y_1$ with $X_1 \subseteq S$ a set of atoms and $Y_1 \subseteq S$ a set of near-litters. We may obviously further stipulate that each element of Y_1 has large intersection with $\bigcup Y'$: if an element of Y_1 does not meet $\bigcup Y'$, we may omit it; if it has small intersection with $\bigcup Y'$, we can omit it and add the elements of the small intersection to X_1 : Y_1 can simply be taken to be the collection of near-litter elements of S with large intersection with $\bigcup Y'$.

Now suppose that the set $X_1 \cup \bigcup Y_1$ which we just constructed as covering $\bigcup Y'$ contains an atom z not in S (and so not in X' , and not in X_1 , so belonging to some element of Y_1) nor in $\bigcup Y'$. Choose an atom w in the same near-litter in $Y_1 \subseteq S$ to which z belongs, belonging to $\bigcup Y'$ (recalling that all elements of Y_1 meet $\bigcup Y'$) but not belonging to S (and so not to X'). A substitution extension of the locally small bijection exchanging z and w , fixing all other anomalies of elements of S , and fixing all atoms in S and atomic parents of near-litters in S will both move and not move Z by considerations now familiar. We have shown that all anomalies of the near-litters belonging to Y_1 which do not belong to $\bigcup Y'$ belong to S .

It follows that Z is the symmetric difference of a set $X \subseteq S$ of atoms and the union (or the complement of the union) of a set $Y \subseteq S$ of near-litters: Z clearly has small symmetric difference X from either the union or the complement of the union of the set $Y = Y_1$ of all near-litters in S which have large intersection with $\bigcup Y'$, and all elements of the symmetric difference X are elements

of S .

the external size of iterated power sets of clans (analysis of orbits):

It is important to note (to appreciate that the induction is sound) that we are in this analysis only considering symmetries in the sense of Π_α^1 , as these are the only symmetries which enter into defining the set Ω_α which we are concerned to show to be of size μ . These symmetries have nice properties following directly from our induction hypotheses and the fact that Π_α^1 is basically the disjoint union of two isomorphic copies of Π_β .

We argue that $\mathcal{P}_*^n(\text{clan}[B])$ is of size μ for each B , $n \leq |B| + 1$. This is certainly true for $n = 0$. The results above on the extent of symmetric power sets of clans show that this is true for $n = 1$: any clan clearly has exactly μ subsets with small symmetric difference from small or co-small unions of litters.

We recall that we refer to strong supports of elements of a $\mathcal{P}_*^{n+1}(\text{clan}[B])$ satisfying the restriction that their elements belong only to $\text{clan}[C]$'s and $\text{litters}[C]$'s with $C \ll B_n$, guaranteed to exist by inductive hypothesis, as “relevant supports”. We further recall that a relevant support of an element of V_ω (the only kind of object which can belong to more than one iterated power set of a clan) is empty.

For any object x with relevant support S with order $<_S$, notice that $\rho(x)$ will have relevant support $\rho(S)$ with order $\rho(<_S)$. The conditions on a relevant support are invariant under application of an allowable permutation. We can then note that we can define a function $\chi_{x,S}$ such that $\chi_{x,S}(\rho(<_S)) = \rho(x)$ for each allowable permutation ρ . To see that this is true, note that if $\rho(<_S) = \rho'(<_S)$, we must have $\rho(x) = \rho'(x)$, because $\rho' \circ \rho^{-1}$ will fix every element of the support of x . We call the functions $\chi_{x,S}$ “coding functions”: ranges of coding functions are orbits under the allowable permutations.

We argue that there are $< \mu$ coding functions with range included in each $\mathcal{P}^{n+1}(\text{clan}[B])$, where $n \leq |B|$.

We define the complexity of a power set $\mathcal{P}^{n+1}(\text{clan}[B])$ as the minimum element of B_n , or λ if B_n is empty. The complexity of a coding function is defined as the smallest complexity of a $\mathcal{P}^{n+1}(\text{clan}[B])$ which includes its range (there is only one such iterated power set unless the coding

function is one whose sole value is a hereditarily finite pure set). Note that the complexity of an iterated power set of a clan which includes the parent of an element of $\mathbf{litters}[A]$ [other than a parent which happens to be a hereditarily finite pure set], that is, the complexity of a $\mathcal{P}^{|B|-|A|+1}(\mathbf{clan}[B])$, is the minimum element of $B_{|B|-|A|} = A$. In the odd case where the parent is a hereditarily finite pure set, the complexity will be less than or equal to the minimum element of A .

Observe that the domain of a coding function is the orbit under the allowable permutations of a relevant support order $<_S$. We can characterize all such orbits by a stereotyped set of information.

Definition (orbit specification): The *orbit specification* of $<_S$ is defined as the function which takes each γ less than the order type of $<_S$ to a tuple consisting of the following components:

1. The first component is 0 if S_γ is an atom, 1 if S_γ is a near-litter.
2. The second component is the index of the clan of which S_γ is an element or subset.
3. The third component is the index δ of the S_δ of which S_γ (an atom) is an element, if there is one, and otherwise is κ .
4. The fourth component, in case S_γ is a near-litter and the parent of S_γ is a set, is the coding function g_γ such that the parent of S_γ is $g_\gamma([<_S]_\gamma)$, where $[<_S]_\gamma$ is the restriction of the initial segment of $<_S$ of order type γ to clans $\mathbf{clan}[C]$ and litter sets $\mathbf{litters}[C]$ for $C \leq B_n$, where the parent of S_γ belongs to $\mathcal{P}^{n+1}(\mathbf{clan}[B])$ [empty if the parent of S_γ is a hereditarily finite pure set], and otherwise is 1.
5. The fifth component, in case S_γ is a near-litter and the parent of S_γ is an atom, is the index δ of the atom S_δ which is the parent of S_γ , if there is one, and otherwise is κ .

Note that a coding function appearing in the specification of the orbit of an order $<_S$ which is a relevant strong support for an element of a set $\mathcal{P}^{n+1}(\mathbf{clan}[A])$ where $n \leq |A|$ will either be of lower complexity than $\mathcal{P}^{n+1}(\mathbf{clan}[A])$ or will be of the same complexity but with domain elements of order type smaller than the order type of $<_S$.

We need to establish that orbit specifications indeed specify orbits in orders on support sets. It should be clear that if ρ is an allowable permutation, $\rho(<_S)$ has the same orbit specification as $<_S$. The considerations in point 6a under the inductive hypotheses can be used to show that if $<_S$ and $<_T$ have the same orbit specification, there is an allowable permutation ρ such that $\rho(<_S) = <_T$. The only interesting case is the case in which S_γ and T_γ are near-litters and have set parents: their respective parents are then of the form $g_\gamma([<_S]_\gamma)$ and $g_\gamma([<_T]_\gamma)$, where g_γ is a coding function – so if we have already defined a locally small bijection with substitution extension sending each S_δ to T_δ for $\delta < \gamma$ and so sending $[<_S]_\gamma$ to $[<_T]_\gamma$, its substitution extensions will send the parent of S_γ to the parent of T_γ , and we can then arrange to send S_γ to T_γ as described in the proof of point 6a under the inductive hypotheses. We have indicated the verification that orbit specifications indeed specify orbits.

We now present the argument for limited size of sets of coding functions. The goal to be proved is that there are $< \mu$ coding functions with range $\mathcal{P}^{n+1}(\text{clan}[B])$ acting on relevant supports whose associated orders have any fixed order type $\gamma < \kappa$ on the hypothesis that there are $< \mu$ coding functions with any range which either have smaller complexity or have the same complexity but act on support orders of length less than γ [once we have shown this we will have shown that there are $< \mu$ coding functions with this range independently of the order type of their domain elements] .

On the inductive hypotheses, there will be $< \mu$ specifications for orbits of support orders of length $< \gamma$ which can be relevant supports of elements of $\mathcal{P}^{n+1}(\text{clan}[B])$, because all coding functions appearing in such specifications will either be of lower complexity than that of $\mathcal{P}^{n+1}(\text{clan}[B])$ or of the same complexity but having domain elements of length $< \gamma$, and orbit specifications are otherwise rather small objects (lists of data of length $< \kappa$).

If $n = 0$, the value of any such coding function will be determined as the symmetric difference of atoms in certain positions in the support order input and either the union of near-litters at certain positions in the support order input or the complement of the union of near-litters at certain support order positions, by the Clan Subset Support Lemma. For each of $< \mu$ possible specifications of the orbit in which the support

order input lies, we have no more than 2^κ possible coding functions, for a total of $< \mu$ coding functions in this case.

It remains to consider the case $n > 0$.

We demonstrate that a coding function χ is completely determined by a specification of the orbit which is its domain and a set of coding functions of lower complexity: let x be an element of the range $\mathcal{P}_*^{n+1}(\text{clan}[B])$ of χ (where $0 < n \leq |B|$), with $x = \chi(<_S)$ (so of course S is a relevant support for x). For each $y \in x$, choose a relevant strong support T so that T end extends the appropriate restriction of $<_S$ (remove from $<_S$ those items not taken from a $\text{clan}[C]$ or $\text{litters}[C]$ with C equal to or downward extending B_{n-1} ; the order $<_T$ on the support T chosen for a $y \in x$ will be an end extension of this restriction). This yields a set of coding functions for elements y of x , all of complexity the minimum of B_{n-1} , so less than the complexity of χ , the minimum of B_n .

We claim that this set of coding functions along with $<_S$ determines x and so χ exactly: we claim that x is exactly the set of all $\chi_{y,T}(<_{T'})$ where $<_{T'}$ end extends the appropriate restriction of $<_S$. Every element of x is of the form $\chi_{y,T}(<_T)$, of course: but further, $\chi_{y,T}(<_{T'})$ belongs to x , too, because we can construct (by the methods of point 6a under the inductive hypotheses) a locally small bijection which adjusts $<_T$ to $<_{T'}$ and in addition fixes all elements of the domain of $<_S$ (noting that $<_S \cup <_T$ and $<_S \cup <_{T'}$ can be extended to orders on strong supports with the correct relationship to one another), and an allowable permutation extending this will send $y = \chi_{y,T}(<_T)$ which is in x to $\chi_{y,T}(<_{T'})$, and will fix x by support considerations, so $\chi_{y,T}(<_{T'}) \in x$ as well. And further, this procedure will work to compute the value of χ at any $<_{S'}$ in its domain, since everything in sight commutes with uniform application of an allowable permutation: so χ is exactly specified by the orbit of $<_S$ and the collection of functions $\chi_{y,T}$.

Now a coding function $\chi_{x,S}$ acting on support orders of length γ and with range in $\mathcal{P}^{n+1}(\text{clan}[B])$ ($n > 0$) is seen to be determined by one of $< \mu$ possible specifications for the orbit in support orders which is its domain, and a subset of the set of coding functions of lower complexity (the set of coding functions of lower complexity is of cardinality $< \mu$; and this set has $< \mu$ subsets because μ is strong limit) so there are $< \mu$

such coding functions, and further it follows immediately that there are $< \mu$ coding functions with this range.

The application of $< \mu$ coding functions with range $\mathcal{P}^{n+1}(\text{clan}[B])$ to μ orders on strong supports will generate no more than μ elements in the iterated symmetric power set $\mathcal{P}_*^{n+1}(\text{clan}[B])$, and this is sufficient to see that the set Ω_α is of cardinality no more than μ (and in fact exactly μ because it includes a clan).

verification of the extension property and induction hypothesis 1:

We demonstrate the extension property for Π_α^1 . Suppose that ρ_0 is a locally small bijection acting on atoms in the closure of Π_α^1 . Let ρ_1 be the restriction of ρ_0 to atoms in the transitive closure of Π_β and let ρ_2 be the restriction of ρ_0 to the transitive closure of $\Pi_\alpha^1 \setminus \Pi_\beta$. Let ρ'_0 be a substitution extension of ρ_1 in the sense of Π_β , and let ρ'_2 be a substitution extension of $\sigma_\beta^{-1} \rho_2 \sigma_\beta$ in the sense of Π_β . The map $\rho'_1 \cup \sigma_\beta \rho'_2 \sigma_\beta^{-1}$ is readily seen to be a substitution extension of ρ_0 . Here we are exploiting the fact that Π_α^1 is in effect the disjoint union of two copies of Π_β .

We state the conditions relating allowability of a permutation in the sense of Π_α^1 and in the sense of Π_α . If ρ is allowable in the sense of Π_α , it is readily extended to **junk** in such a way as to be allowable in the sense of Π_α^1 : for x in **junk**, define $\rho(x)$ as $h^{-1}(\rho(h(x)))$. We will refer to this result informally in the terms “a permutation allowable in the sense of Π_α is allowable in the sense of Π_α^1 ”. A permutation ρ which is allowable in the sense of Π_α^1 is allowable in the sense of Π_α [after disregarding its values on **junk**, which is not in the transitive closure of Π_α] iff it satisfies the condition $\rho(h(x)) = h(\rho(x))$ [which is of course a restrictive condition on how it acts on **junk**]. The issue in both cases is that the difference between Π_α^1 and Π_α has precisely to do with replacing the parent $x \in \text{junk}$ of each litter included in $\text{clan}[\{\beta\}]$ with $h(x) \in \Omega_\alpha$.

We prove an

Extension Lemma: Let ρ_0 be an allowable permutation in the sense of Π_β (considered as implemented as a map on the atoms in the transitive closure of Π_β) and let ρ_1 be a locally small bijection on atoms in the transitive closure of $\Pi_\alpha \setminus \Pi_\beta$, agreeing with ρ_0 where their domains intersect (which will be in $\text{parents}[\emptyset]$ and $\text{clan}[\emptyset]$)

only). Then there is an allowable permutation ρ which extends ρ_0 , extends ρ_1 , and has no exceptions other than exceptions of ρ_0 and elements of the domain of ρ_1 .

Proof of Extension Lemma: Notice that ρ_1 has no action on **junk**.

We will define a map ρ_2 on **junk**.

Fix maps $\rho_{L,M}$, bijections from $L \setminus \text{dom}(\rho_1)$ to $M \setminus \text{dom}(\rho_1)$, for each pair of litters L, M in the same clan with index whose maximum is β .

For each x such that $h(x) \in \text{clan}[\emptyset]$, we define $\rho_2(x)$ as $h^{-1}(\rho_0(h(x)))$.

We define all other values of ρ_2 by recursion along the order $<_h$ used in the construction of h above. Suppose that x is the $<_h$ -least element of **junk** at which ρ_2 is not yet defined. Let ρ_x be a substitution extension in the sense of Π_α^1 of the union of the part of ρ_2 already defined and the locally small bijection ρ_1 , also agreeing with each map $\rho_{L,\rho_x(L)^\circ}$: the construction of such a map ρ_x is described under point 6b under the inductive hypotheses. The value $\rho_x(h(x))$ is uniquely determined (it is not affected by the arbitrary assignment of values of ρ_x to elements of **junk** for which ρ_2 has not yet been defined), because $h(x)$ has a strong support in which any irregular atoms that occur as parents of litters already have ρ_2 defined, by the construction of h . We define $\rho_2(x)$ as $h^{-1}(\rho_x(h(x)))$. We iterate this until ρ_2 is defined on all of **junk**, then take ρ to be the union of ρ_1 and the unique (on atoms in the transitive closure of $\Pi_\alpha \setminus \Pi_\beta$) substitution extension of $\rho_2 \cup \rho_1$ which extends each $\rho_{L,\rho(L)^\circ}$. If we disregard the action of this ρ on **junk**, it satisfies the conditions to be allowable on Π_α , and it clearly has only the exceptions indicated.

We can use the Extension Lemma to verify the extension property in the sense of Π_α : if ρ_0 is a locally small bijection in the sense of Π_α , let ρ_1 be a substitution extension in the sense of Π_β of the restriction of ρ_0 to the transitive closure of Π_β , and let ρ_2 be the restriction of ρ_0 to the transitive closure of $\Pi_\alpha \setminus \Pi_\beta$. The maps ρ_1 and ρ_2 satisfy the hypotheses of the Extension Lemma, and the extension of these maps obtained will be a substitution extension of the original ρ_0 .

We can use the Extension Lemma to verify induction hypothesis 1. The action of an allowable permutation in the sense of Π_α on the atoms in

the transitive closure of Π_β is allowable in the sense of Π_β . To show that the symmetry on iterated power sets of clans whose index has maximum $< \beta$ is the same in the senses of the two parent functions, we also need to show that any allowable permutation in the sense of Π_β extends to an allowable permutation in the sense of Π_α . This follows from the Extension Lemma: take ρ_0 to be any allowable permutation in the sense of Π_β , take ρ_1 to be empty, and apply the Extension Lemma to get the desired extension.

Finally we comment that the extension property can be seen to hold at limit stages as well: let ρ_0 be a locally small bijection in the sense of a Π_λ whose domain includes all irregular atoms. We can extend successively larger restrictions of ρ_0 to successive Π_α 's for $\alpha < \lambda$, applying the successor stage construction at successors and taking unions at limits (we can use the extension lemma to ensure that the successive approximations that we construct are nested), until the extension of the full ρ_0 to Π_λ is obtained. Similar considerations apply to induction hypothesis 1.

the need to preserve notions of symmetry: To establish that the desired formula for parent sets continues to hold at $A = \{\beta\}$ we need to show that the iterated symmetric power sets appearing in the definition of the complex parent set have the same extension when Π_α is used to define allowable permutations in place of Π_α^1 : we will argue that the definition of $\mathcal{P}_*^n(\text{clan}[B])$ for $n \leq |B|, B \ll \{\beta\}$ using Π_α as the parent function agrees with the definition of $\mathcal{P}_*^n(\text{clan}[B])$ for $n \leq |B|, B \ll \{\beta\}$ using Π_α^1 as the parent function. This is also needed to verify the induction hypothesis on existence of external isomorphisms, and to verify the inductive hypothesis concerning the existence of relevant supports.

the argument for preservation of notions of symmetry: We show by induction on n that $\mathcal{P}_*^n(\text{clan}[B])$ has the same extension for each n with $1 \leq n \leq |B|$ in terms of Π_α as in terms of Π_α^1 .

basis step: This result for $n = 1$ follows from the results on the extent of symmetric power sets of clans: these are uniquely specified in a way quite independent of the details of the parent function used.

induction step: We now argue that if $\mathcal{P}_*^n(\text{clan}[B])$ has the same extension in terms of both Π_α^1 and Π_α , so does $\mathcal{P}_*^{n+1}(\text{clan}[B])$, where $1 \leq n < |B|$ and the maximum element of B is β (where the maximum element is not β there is no issue, as Π_α symmetry is evidently the same as Π_β symmetry in this case). One direction of this is easy, as all Π_α -allowable permutations are also Π_α^1 -allowable permutations, so any object c invariant under all Π_α^1 -allowable permutations fixing all elements of a support set S is also invariant under all Π_α -allowable permutations fixing all elements of the same support set. So what we actually need to show is that a subset X of $\mathcal{P}_*^n(\text{clan}[B])$ with a strong support S relative to Π_α -allowable permutations also has support S relative to Π_α^1 permutations.

Strong supports exist in the sense of the parent function Π_α because of the way h is constructed, which ensures that each of the elements of $\text{parents}[\{\beta\}]$ (in the sense of Π_α , that is, Ω_α) has Π_α -strong support computable by recursion along $<_h$ from its Π_α^1 -strong support, computed from the Π_β -strong support of its preimage under σ_β , and the Π_α -strong supports computed for elements of $\text{parents}[\{\beta\}] = \Omega_\alpha$ considered earlier in the construction of h .

Let ρ be a Π_α^1 -allowable permutation fixing each element of S , a Π_α -strong support of $X \subseteq \mathcal{P}_*^n(\text{clan}[B])$. Let c belong to X . The element c has a Π_α^1 -relevant strong support T by inductive hypothesis, which is also a Π_α -strong support for c (to see this, observe that any element of T belongs to $\text{clan}[C]$ or $\text{litters}[C]$ with $C \ll B_{n-1}$, which has at least two elements: the only way T could fail to be Π_α -strong would be if it contained an element of $\text{litters}[\{\beta\}]$). Now extend T to a Π_α -strong support U including S , with associated well-ordering $<_U$.

Construct a locally small bijection ρ_0 sending each atom which either belongs to U or is the parent of a near-litter belonging to U to its image under ρ , and further ensure that each anomaly of an element of U or preimage under ρ of an anomaly of an element of $\rho(U)$ and each exception of ρ has the same image under ρ_0 that it has under ρ , and that ρ_0 respects each near-litter element u of U in the sense that ρ_0 maps elements of u to elements of

$\rho(u)$ and non-elements of u to non-elements of $\rho(u)$. Choose a Π_α -allowable substitution extension ρ' of ρ_0 : ρ' agrees with ρ on U , at atoms by construction and at near-litters for reasons which by now should be familiar: if not, consider the first near-litter $u \in U$ such that $\rho(u) \neq \rho'(u)$; the parents of these two near litters are equal because ρ' is Π_α^1 -allowable and has the same action on a support of the parent as ρ ; ρ' must then map some element of u to a non-element of $\rho(u)$ or vice versa, but we have arranged that this element is forced to be an exception of ρ' (by forcing ρ to agree with ρ' at anomalies of relevant near-litters), and ρ' has no exceptions at which it disagrees with ρ . Thus ρ sends each element of U to its image under ρ , and thus fixes X because it fixes all elements of S , and satisfies $\rho'(c) = \rho(c)$ because ρ and ρ' have the same values on T , a Π_α -support for c . Thus $\rho'(c) = \rho(c) \in X$: $\rho^{-1}(c) \in X$ by an identical argument, so ρ fixes X , so S is also a Π_α^1 -support for X .

preserving external isomorphism conditions and support restrictions:

Note that we will enforce our inductive hypothesis about existence of external isomorphisms between initial segments of natural models of type theory if the meaning of \mathcal{P}_* in descriptions of these parent sets is preserved in our passage from Π_α^1 to Π_α : σ_β will witness all these isomorphisms: the default natural model with top type $\mathcal{P}_*^{n+1}(\text{clan}[A])$ and bottom type $\text{clan}[A]$ ($0 < n < |A|$) will be an isomorphic copy of the model with top type $\mathcal{P}_*^{n+1}(\text{clan}[A \setminus A_n])$ and bottom type $\text{clan}[A \setminus A_n]$ by application of the isomorphism witnessed by σ_β and a finite series of external isomorphisms found at earlier stages, and application of the result that notions of symmetry are preserved.

We need to verify that we have a strong support for any element of a $\mathcal{P}_*^{n+1}(\text{clan}[B])$ involving clans and near-litters taken from clans with index equal to or downward extending B_n (a relevant support). For B not containing β , this is direct. For B containing β , where B_n is not empty, an element of $\mathcal{P}_*^{n+1}(\text{clan}[B])$ has a relevant support relative to Π_α^1 , copied from the relevant support of the analogous object in $\mathcal{P}_*^{n+1}(\text{clan}[B \setminus \{\beta\}])$, which is still a support of it relative to Π_α because the allowable permutations with respect to Π_α are a subset of those allowable with respect to Π_α^1 . This does depend on the result shown above that $\mathcal{P}_*^{n+1}(\text{clan}[B])$ has no new elements with respect to Π_α

(that notions of symmetry are preserved). That we have a relevant support for a $\mathcal{P}_*^{n+1}(\mathbf{clan}[B])$ where $\beta \in B$ and B_n is empty is evident, because relevance is no constraint in this case.

This completes the proof that the construction of Π (as Π_λ) works. Slightly more is needed to complete the main argument.

cardinalities of subsets of clans: We argue that a symmetric bijection between subsets of a clan has small symmetric difference from the identity map on its domain. Suppose that f is a bijection from $X \subseteq \mathbf{clan}[A]$ to $Y \subseteq \mathbf{clan}[A]$ with extended strong support S and with large symmetric difference from the identity, so there is a large subset X' of its domain on which it is not fixed. By earlier results, X' must have large intersection with a litter L . Choose elements a, b of L such that $a, b, f(a), f(b)$ are all distinct and none of them belong to S (consider the fact that orbits in f are small sets). A substitution extension of the locally small bijection which exchanges a, b and fixes $f(a), f(b)$ and all atoms in S is seen both to fix X and to move it. The substitution extension is seen to fix near-litter elements of S by considering the first litter in $<_S$ which it moves as in arguments above, noticing that it must contain an exception or an image of an exception of the substitution extension. . . which has no exceptions or images of exceptions which belong to litters in S .

This tells us that the litters have distinct κ -amorphous cardinalities in the FM interpretation (a κ -amorphous set being one which has only small and co-small subsets, and a κ -amorphous cardinal being the cardinality of such a set).

an important injection: There is a symmetric injection from $\mathcal{P}_*(\mathbf{parents}[A])$ to $\mathcal{P}_*^2(\mathbf{clan}[A])$, associating each symmetric subset X of $\mathbf{parents}[A]$ with the set union of the collection of local cardinals with parents in X . This witnesses the important inequality

$$|\mathcal{P}_*^2(\mathbf{clan}[A])| \geq |\mathcal{P}_*(\mathbf{parents}[A])|$$

holding in the FM interpretation. We do not actually need the result about cardinality just above to show this, as this map is obviously symmetric and an injection. The preceding result is a nice result, though, and justifies the use of the term “local cardinal”.

motivational remark: This situation is a major aim of the elaborate machinery of our construction. The FM interpretation is designed so that the power set of a clan is almost amorphous, and its structure reveals nothing about the structure of the parent set that the FM interpretation can see, but the double power set of a clan contains structure which the FM interpretation can see as parallel to the structure of the power set of the parent set (not just the parent set itself!). This is part of what enables us to fit the cardinalities of iterated power sets of clans together into the unlikely structure of a tangled web.

The main theorem: We can now prove that the map τ on nonempty clan indices defined by the equation $\tau(A) = |\mathcal{P}_*^2(\text{clan}[A])|$ is a tangled web, from which the main result of the paper that NF is consistent follows at once.

definition and verification of the tangled web: All cardinality facts mentioned here are those of the FM interpretation.

Define $\exp(|X|)$ as $|\mathcal{P}_*(X)|$.

We need to verify that $\exp(\tau(A)) = \tau(A_1)$ if $|A| \geq 2$. This is equivalent to showing that $|\mathcal{P}_*^3(\text{clan}[A])| = |\mathcal{P}_*^2(\text{clan}[A_1])|$ [whence it is straightforward to show that $|\mathcal{P}_*^{n+2}(\text{clan}[A])| = |\mathcal{P}_*^2(\text{clan}[A_n])|$ when $|A| > n$]. We have from the inequality witnessed by the injection described just above and the formula for parent sets that

$$\begin{aligned} |\mathcal{P}_*^2(\text{clan}[A_1])| &\geq |\mathcal{P}_*(\text{parents}[A_1])| \\ &\geq |\mathcal{P}_*(\mathcal{P}_*^{|A|-|A_1|+1}(\text{clan}[A]))| = |\mathcal{P}_*^3(\text{clan}[A])|. \end{aligned}$$

On the other hand

$$|\mathcal{P}_*^3(\text{clan}[A])| \geq |\mathcal{P}_*^2(\text{parents}[A])| \geq |\mathcal{P}_*^2(\text{clan}[A_1])|.$$

This verifies the naturality property of tangled webs.

The elementarity property of tangled webs falls directly out of the construction. We need to show that the first order theory of a natural model (all natural models discussed here being those of the FM interpretation) of TST_n whose base type has cardinality $\tau(A)$ depends only on $A \setminus A_n$, the set of the smallest n elements of A , where $|A| > n$. This reduces to consideration of default natural models of TST_n whose base

type is $\mathcal{P}_*^2(\mathbf{clan}[A])$ and whose top type is $\mathcal{P}_*^{n+1}(\mathbf{clan}[A])$. This model is the image under the action of a bijection on atoms [in the ground model of ZFA] of the default natural model of type theory whose base type is $\mathcal{P}_*^2(\mathbf{clan}[A \setminus A_n])$ and whose top type is $\mathcal{P}_*^{n+1}(\mathbf{clan}[A \setminus A_n])$, by the construction (the construction actually provides an external isomorphism between the default natural models of \mathbf{TST}_{n+2} with base types $\mathbf{clan}[A]$ and $\mathbf{clan}[A \setminus A_n]$, a composition of maps σ_β): these models have the same first-order theory because they are isomorphic models of the appropriate initial segment of type theory from the standpoint of the ground model of ZFA, so the theory of the model considered initially depends only on $A \setminus A_n$.

At this point the main result of the paper (the consistency of NF) is proved.

a possible simplification of the formula for parent sets: If we define

$B \ll_1 A$ as $B \ll A \wedge |B| - |A| = 1$ (in other words, $B = A \cup \{\beta\}$ for some β dominated by all elements of A), we could carry out the proof using the formula (for nonempty \mathbf{clan} indices A)

$$\mathbf{parents}[A] = \mathbf{clan}[A_1] \cup \bigcup_{B \ll_1 A} \mathcal{P}_*^2(\mathbf{clan}[B]).$$

or equivalently

$$\mathbf{parents}[A] = \mathbf{clan}[A_1] \cup \bigcup_{\beta < \mathbf{min}(A)} \mathcal{P}_*^2(\mathbf{clan}[A \cup \{\beta\}]).$$

One can see from the proof just above that this is sufficient to establish the naturality property of the tangled web. We have decided not to adopt this “simplified” form basically because nothing much is gained in terms of complexity of the proof: it is *still* necessary to show preservation of notions of symmetry for all $\mathcal{P}_*^n(\mathbf{clan}[B])$ with $n \leq |B|$ (this n does not reduce to 2) because the preservation of symmetry is needed to verify the isomorphism conditions as well as the parent set formula. It is also the case that this is only nominally less evilly tangled than the original formula.

an outline of an interpretation of tangled type theory: It should be noted that the external isomorphisms between iterated power sets in the

natural models based on the tangled web are such that various iterated power sets can in fact be identified in such a way as to produce a model of the tangled type theory TTT_λ . This is important in connection with application of the results indicated in section 6.1, which allow us to draw the further conclusion that there is an ω -model of NF. We briefly indicate how to do this. Our aims here are restricted to a compact description of the interpretation and a general description of the reasons why it is an interpretation: we feel free to do this as our main result does not depend on this; the reason that we discuss it is that it makes it easier for us to give an indication of reasons why the existence of an ω -model of NF follows from our construction, which is important for corollaries mentioned in the conclusions section below.

Type α for each $\alpha < \lambda$ is conveniently implemented as $\mathcal{P}_*^2(\text{clan}(\{\alpha\}))$. A scheme of bijections $E_{A,n} : \mathcal{P}_*^2(\text{clan}([\{\alpha\}])) \rightarrow \mathcal{P}_*^{n+2}(\text{clan}[A])$ are presented, for each A for which $\min(A_n) = \alpha$. These bijections implement a scheme of identification of each $\mathcal{P}_*^{n+2}(\text{clan}[A])$ for which $\min(A_n) = \alpha$ with type α of the interpreted tangled type theory.

The embedding $E_{\{\alpha\},0}$ is of course the identity map.

The embedding $E_{\{\alpha,\beta\},1}$ (where $\alpha > \beta$) is in each case a hereditarily symmetric bijection from $\mathcal{P}_*^2(\text{clan}[\{\alpha\}])$ to $\mathcal{P}_*^3(\text{clan}[\{\alpha,\beta\}])$: we have just shown that such bijections exist.

If the embedding E_A has been defined, and δ dominates all members of A , we define $E_{\{\delta\} \cup A,n}(x)$ as $\sigma_\delta(E_{A,n}(x))$, noting that E_A is already known to map $\mathcal{P}_*^2(\text{clan}(\{\alpha\}))$ to $\mathcal{P}_*^{n+2}(\text{clan}[A])$, while σ_δ [defined in the construction above] acts as an isomorphism between natural models of initial segments of type theory with clans as base types which (among other things) sends $\mathcal{P}_*^{n+2}(\text{clan}[A])$ to $\mathcal{P}_*^{n+2}(\text{clan}[\{\delta\} \cup A])$.

If $E_{A,n}$ is defined and each element of A dominates δ , $E_{A \cup \{\delta\},n+1}$ is intended to map $\mathcal{P}_*^2(\text{clan}[\{\min(A_n)\}])$ to $\mathcal{P}_*^{n+3}(\text{clan}[A \cup \{\delta\}])$. Let α denote $\min(A_n)$ and let β denote $\min(A_{n-1})$. We may suppose that we have already defined the map $E_{\{\alpha,\beta\},1}$ from $\mathcal{P}_*^2(\text{clan}(\{\alpha\}))$ to $\mathcal{P}_*^3(\text{clan}[\{\alpha,\beta\}])$ and the map $E_{A \cup \{\delta\},n}$ from $\mathcal{P}_*^2(\text{clan}[\{\beta\}])$ to $\mathcal{P}_*^{n+2}(\text{clan}[A \cup \{\delta\}])$. Define $E_{A \cup \{\delta\},n+1}(x)$ as $E_{A \cup \{\delta\},n}(\sigma_\alpha^{-1}(E_{\{\alpha,\beta\},1}(x)))$. The application of σ_α^{-1} converts elements of $\mathcal{P}_*^3(\text{clan}[\{\alpha,\beta\}])$ to elements of $\mathcal{P}_*^3(\text{clan}[\{\beta\}])$ by the application of an external isomorphism (acting on atoms) of a relevant natural model of an initial segment of

type theory.

The crucial feature is that that for any x, y , if $E_{A,n}(x)$ and $E_{A,n+1}(y)$ are defined and $E_{A',n'}(x)$ and $E_{A',n'+1}(y)$ are defined, then $E_{A,n}(x) \in E_{A,n+1}(y) \leftrightarrow E_{A',n'}(x) \in E_{A',n'+1}$: types identified via the scheme of bijections agree about membership facts.

For $\beta < \alpha < \lambda$, the membership $x \in_{\beta,\alpha} y$ of the interpretation of tangled type theory whose construction we outline here is defined as holding, for $x \in \mathcal{P}_*^2(\mathbf{clan}[\{\beta\}])$ and $y \in \mathcal{P}_*^2(\mathbf{clan}[\{\alpha\}])$, just in case $E_{A,n}(x) \in E_{A,n+1}(y)$ for some (and so for any) A such that $\min(A_n) = \alpha, \min(A_{n+1}) = \beta$: concretely, $x \in_{\beta,\alpha} y$ iff $E_{\{\alpha,\beta\},0}(x) \in E_{\{\alpha,\beta\},1}(y)$, that is, $\sigma_\alpha(x) \in E_{\{\alpha,\beta\},1}(y)$. The verification of extensionality and comprehension follows in a straightforward manner by considering the identifications of the types with suitable segments of natural models of type theory (natural models in the sense of the FM interpretation, of course).

It is important to note that the relations $\in_{\beta,\alpha}$ of the interpretation of tangled type theory are not set relations in the FM interpretation, because of the role of the external maps σ_δ derived from the construction of the FM interpretation in the definition of the bijections generating the interpretation of tangled type theory.

9 Conclusions and questions

The conclusions to be drawn about NF are rather unexciting ones.

By choosing the parameter λ to be larger (and so to have stronger partition properties) one can show the consistency of a hierarchy of extensions of NF similar to extensions of NFU known to be consistent: one can replicate Jensen's construction of ω - and α -models of NFU to get ω - and α -models of NF (details given above). One can show the consistency of NF + Rosser's Axiom of Counting (see [13]), Henson's Axiom of Cantorian Sets (see [4]), or the author's axioms of Small and Large Ordinals (see [6], [7], [15]) in basically the same way as in NFU.

It seems clear that this argument, suitably refined, shows that the consistency strength of NF is exactly the minimum possible on previous information, that of TST + Infinity, or Mac Lane set theory (Zermelo set theory with comprehension restricted to bounded formulas). Actually showing that the consistency strength is the very lowest possible might be technically tricky, of course. We have not been concerned to do this here. It is clear from what is done here that NF is much weaker than ZFC.

By choosing the parameter κ to be large enough, one can get local versions of Choice for sets as large as desired, using the fact that any small subset of a type of the structure is symmetric. The minimum value ω_1 for κ already enforces Denumerable Choice (Rosser's assumption in his book) or Dependent Choices. It is unclear whether one can get a linear order on the universe or the Prime Ideal Theorem: that would require major changes in this construction. But certainly the question of whether NF has interesting consequences for familiar mathematical structures such as the continuum is answered in the negative: set κ large enough and what our model of NF will say about such a structure will be entirely in accordance with what our original model of ZFC said. It is worth noting that the models of NF that we obtain are not κ -complete in the sense of containing every subset of their domains of size κ ; it is well-known that a model of NF cannot contain all *countable* subsets of its domain. But the models of TST from which its theory is constructed will be κ -complete, so combinatorial consequences of κ -completeness expressible in stratified terms will hold in the model of NF (which could further be made a κ -model by making λ large enough).

The question of Maurice Boffa as to whether there is an ω -model of TNT (the theory of negative types, that is TST with all integers as types, proposed by Hao Wang ([18])) is settled: an ω -model of NF yields an ω -model of TNT

instantly. This work does not answer the question, very interesting to the author, of whether there is a model of TNT in which every set is symmetric under permutations of some lower type.

The question of the possibility of cardinals of infinite Specker rank (at least in ZFA) is answered, and we see that the existence of such cardinals doesn't require much consistency strength. For those not familiar with this question, the Specker tree of a cardinal is the tree with that cardinal at the top and the children of each node (a cardinal) being its preimages under $\alpha \mapsto 2^\alpha$. It is a theorem of Forster (a corollary of a well known theorem of Sierpinski) that the Specker tree of a cardinal is well-founded (see [2], p. 48), so has an ordinal rank, which we call the Specker rank of the cardinal. NF + Rosser's Axiom of Counting proves that the Specker rank of the cardinality of the universe is infinite; it was unknown until this point whether the existence of a cardinal of infinite Specker rank was consistent with any set theory in which we had confidence. The possibility of a cardinal of infinite Specker rank in ZFA is established by the construction here; we are confident that standard methods of transfer of results obtained from FM constructions in ZFA to ZF will apply to show that cardinals of infinite Specker rank are possible in ZF.

This work does not answer the question as to whether NF proves the existence of infinitely many infinite cardinals (discussed in [2], p. 52). A model with only finitely many infinite cardinals would have to be constructed in a totally different way. We conjecture on the basis of our work here that NF probably does prove the existence of infinitely many infinite cardinals, though without knowing what a proof will look like.

A natural general question which arises is, to what extent are *all* models of NF like the ones indirectly shown to exist here? Do any of the features of this construction reflect facts about the universe of NF which we have not yet proved as theorems, or are there quite different models of NF as well?

10 References and Index

References

- [1] Maurice Boffa, “Entre NF et NFU ,” *Comptes Rendues de l’Academie des Sciences de Paris*, series A, 277 (1973), pp. 821–2.
- [2] Forster, T.E. [1995] Set Theory with a Universal Set, exploring an untyped Universe Second edition. Oxford Logic Guides, Oxford University Press, Clarendon Press, Oxford.
- [3] Hailperin, T. “A set of axioms for logic”. *Journal of Symbolic Logic* 9 (1944), pp. 1-19.
- [4] Henson, C.W. [1973a] Type-raising operations in NF . *Journal of Symbolic Logic* 38 , pp. 59-68.
- [5] Holmes, M.R. “The equivalence of NF -style set theories with ”tangled” type theories; the construction of omega-models of predicative NF (and more)”. *Journal of Symbolic Logic* 60 (1995), pp. 178-189.
- [6] Holmes, M. R. [1998] Elementary set theory with a universal set. volume 10 of the Cahiers du Centre de logique, Academia, Louvain-la-Neuve (Belgium), 241 pages, ISBN 2-87209-488-1. See here for an on-line errata slip. By permission of the publishers, a corrected text is published online; an official second edition will appear online eventually.
- [7] Holmes, M. R. [2001] Strong Axioms of infinity in NFU . *Journal of Symbolic Logic*, 66, no. 1, pp. 87-116.
 (“Errata in ‘Strong Axioms of Infinity in NFU ’ ”, *JSL*, vol. 66, no. 4 (December 2001), p. 1974, reports some errata and provides corrections).
- [8] Jech, Thomas, *The axiom of choice*, North-Holland 1973, p. 46
- [9] Jech, Thomas, *Set theory*, Academic Press 1978, pp. 199-201.
- [10] Jensen, R.B. “On the consistency of a slight(?) modification of Quine’s NF ”. *Synthese* 19 (1969), pp. 250-263.
- [11] Quine, W.V., “New Foundations for Mathematical Logic”. *American Mathematical Monthly* 44 (1937), pp. 70-80.

- [12] Ramsey, F.P., 1926, “The foundations of mathematics, *Proceedings of the London Mathematical Society*, s225 (1), 338384.
- [13] Rosser, J. B. [1978] *Logic for mathematicians*, second edition. Chelsea Publishing.
- [14] Scott, Dana, “Definitions by abstraction in axiomatic set theory”, *Bull. Amer. Math. Soc.*, vol. 61, p. 442, 1955.
- [15] Solovay, R, “The consistency strength of NFUB”, preprint on [arXiv.org](https://arxiv.org/abs/math/9707207), [arXiv:math/9707207](https://arxiv.org/abs/math/9707207) [math.LO]
- [16] Specker, E.P. “The axiom of choice in Quine’s new foundations for mathematical logic”. *Proceedings of the National Academy of Sciences of the USA* 39 (1953), pp. 972-975.
- [17] Specker, E.P. [1962] “Typical ambiguity”. *Logic, methodology and philosophy of science*, ed. E. Nagel, Stanford University Press, pp. 116-123.
- [18] Wang, H. [1952] Negative types.
- [19] Wang, Hao, *Logic, Computers and Sets*, Chelsea 1970, pp. 402-415.
- [20] Whitehead, A. N., and Russell B., *Principia Mathematica*, second edition, 3 vols. Cambridge University Press, 1927.
- [21] Wiener, N., 1914, “A simplification of the logic of relations, *Proceedings of the Cambridge Philosophical Society*, 17: 387390.